Self-Financed Wagering Mechanisms for Forecasting*

Nicolas Lambert
Stanford University

John Langford
Yahoo! Research

Jennifer Wortman
University of Pennsylvania

Yiling Chen
Yahoo! Research

Daniel Reeves
Yahoo! Research

Yoav Shoham
Stanford University

David M. Pennock
Yahoo! Research

ABSTRACT

We examine a class of wagering mechanisms designed to elicit truthful predictions from a group of people without requiring any outside subsidy. We propose a number of desirable properties for wagering mechanisms, identifying one mechanism—weighted-score wagering—that satisfies all of the properties. Moreover, we show that a single-parameter generalization of weighted-score wagering is the only mechanism that satisfies these properties. We explore some variants of the core mechanism based on practical considerations.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Theory

Keywords

Group forecasting, mechanism design, prediction markets

1. INTRODUCTION

Consider a group of people with different estimates about an uncertain variable, for example different probability estimates or different quantile estimates. We seek mechanisms to induce all the group members to truthfully reveal their estimates to each other.

If some patron is willing to pay for the group's information, then the patron can use any number of well-known *scoring rule* payment functions that give group members the incentive to truthfully report their estimates [16, 14]. However, if no patron is willing to subsidize the process, as we assume in this paper, then self-financed or budget-balanced mechanisms are needed.

*Part of this work was done while N. Lambert and J. Wortman were at Yahoo! Research, New York.

EC'08, July 8–12, 2008, Chicago, Illinois, USA.

Prediction markets [17] can be run in a self-financed way, for example by employing double auction or parimutuel mechanisms. Such markets induce aggregate estimates [18, 7], but do not induce truthful individual estimates [3]. An informed trader may simply lack the wealth necessary to move the market to express a belief. Similarly, if a trader has a belief that matches the current market price, there is no incentive to participate. Additionally, double auctions are prone to the thin-market problem: with few traders, the volume of transactions is insufficient attract other traders, hindering price discovery.

Kilgour and Gerchak [10] also propose an adaptation of scoring rules for groups that are budget-balanced. However, they require that all group members wager the same amount of money. Recently, Johnstone [8] adapted the Kilgour-Gerchak scoring rules to account for different wagers, but the resulting mechanism is not truthful. Newsfutures.com employs a form of competitive forecasting for continuous random variables that works well in practice, though is not truthful.

In this paper, we take an axiomatic approach, compiling a list of desirable properties in Section 4 that arguably any mechanism should have: budget balance, anonymity, truthfulness, normality, sybilproofness, individual rationality, and monotonicity. In Section 2, we construct a general space of mechanisms called *wagering mechanisms* which contains many known approaches, including call markets and parimutuel markets. In Section 3, we identify the weighted-score mechanism, a generalization of Kilgour-Gerchak scoring rules, adapted to the elicitation of general distribution properties and with weights a function of amount wagered. We prove that weighted-score wagering satisfies all seven properties.

A natural question to ask is whether there are other mechanisms satisfying this set of properties. We prove that the answer is *no*. More precisely, in Section 5 we prove that the mechanisms that satisfy the first five properties are exactly the weighted-score mechanisms, parameterized by the total amount wagered. Then, relaxing some of the properties, in Section 6 we explore variants of the core mechanism based on practical considerations.

2. PRELIMINARIES

2.1 Model

We consider a wagering setting involving a *principal* and a finite set of agents, or *players* $\mathcal{N} = \{1, \ldots, n\}$. The principal is interested in eliciting certain information on a given random experiment, for example a horse race or the election of a political candidate. We denote by Ω the set of possible outcomes. We assume that the outcome of the random experiment is drawn according to some distribution unknown to the principal, but on which players form beliefs. We model information of interest as *distribution properties*, such as event probabilities or mean, defined in Section 2.3.

Players participate in a *wagering mechanism*, formalized in the next subsection. Wagering mechanisms are essentially one-shot games in which players wager some money and make a prediction. After the realization of the experiment, each player receives a payout that is a function of her performance relative to the performance of other players.

We assume that players are risk neutral and seek to maximize their expected payout with respect to their belief. We argue that linear utility is a reasonable assumption when the money at stake remains small enough. We further assume that players choose their wager up front, before deciding on their prediction.[1] We believe that such a behavior is plausible in a context where players have fixed budgets. In particular, this is consistent with the model and findings of Ali [1] in the case of horse-race parimutuel betting. Ali showed that amounts players are willing to bet are typically small and player utility varies approximately linearly in the range of amounts being bet.

2.2 Wagering Mechanisms

Wagering mechanisms, which cover a large class of one-shot mechanisms, allow players to specify a wager as part of their actions. Wagering mechanisms operate in two steps. In the first step, each player announces a report chosen from a certain *set of possible reports* $\mathcal{R}$ and wagers any positive amount of money. Wagers are deposited in a common pot. The set of allowed reports $\mathcal{R}$ may take various forms, for example, in the case of a horse race, players may be asked to specify a winning horse, or may be asked to lay out the probability of winning for each competing horse. In a second step, after the realization of the experiment, the common pot is divided among the players according to their performance, based on their reports and the true outcome of the experiment. When the mechanism is not budget-balanced, players may receive a bonus, or may have to pay a tax or fee. Using the vector notation $\mathbf{x} = (x_1, \ldots, x_n)$, each player i reporting r_i and wagering m_i gets a nonnegative payout $\Pi_i(\mathbf{r}, \mathbf{m}, \omega)$ specified by the mechanism. Payouts depend on reports and wagers, and on the true outcome of the experiment ω. Player i's net profit is thus $\Pi_i(\mathbf{r}, \mathbf{m}, \omega) - m_i$.

We assume that a player who wagers zero gets zero payout, so without loss of generality, we can use the set of natural numbers to represent the set of players $\mathcal{N} = \mathbb{N}$ since nonparticipation is equivalent to a zero bid.[2] The formal definition of wagering mechanism is given below.

Definition 1. A *Wagering Mechanism* is a tuple $(\mathcal{R}, \Omega, \Pi)$ together with a set of players $\mathbb{N}$, where $\mathcal{R}$ is the set of allowed reports, Ω is the outcome space, and $\Pi = (\Pi_i(\mathbf{r}, \mathbf{m}, \omega))_{i\in\mathbb{N}}$ is the vector of payout functions $\Pi_i : \mathcal{R}^{\mathbb{N}} \times [0, +\infty)^{\mathbb{N}} \times \Omega \mapsto [0, +\infty)$, with $\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = 0$ if $m_i = 0$.

Wagering mechanisms include several well-known instances of betting mechanisms, in particular, parimutuel betting markets [1, 15]. In a parimutuel market, players wager on mutually exclusive and exhaustive events $E_1, \ldots, E_m$ (where, for example, E_i = "horse i wins the race"). Players lose their wagers when the true outcome is not what they bet, while winning players share the total money wagered in proportion to their own wager. Such a market is a wagering mechanism with the set of reports $\mathcal{R} = \{E_1, \ldots, E_m\}$, and a payout

$$\Pi_i = \mathbb{1}_{\omega\in r_i} \frac{m_i}{\sum_j m_j \mathbb{1}_{\omega\in r_j}} \sum_j m_j \ ,$$

when player i reports r_i and wagers m_i dollars, and where $\mathbb{1}_{\omega\in A}$ is 1 if $\omega \in A$ and 0 otherwise.

Call markets [9, 13, 4] can also be viewed as wagering mechanisms. In a typical call market for binary events, participants trade on contingent contracts, each contract corresponding to an outcome, paying off \$1 if the outcome becomes true and \$0 otherwise. Traders may submit orders, which indicate the maximum price they are willing to pay for a certain contract and a maximum amount of money to spend. Whenever possible, the market matches orders for one outcome with orders for the opposite one. When the space of outcomes is large, more exotic betting languages can be used to allow, for example, betting on combinations of events. This includes Boolean betting where players bet on Boolean formulas of events [5] and permutation betting where players bet on properties of the final ranking of competing candidates [2]. Such combinatorial betting may use call markets with multilateral order matching to clear bets, which can be modeled as wagering mechanisms. However the payout functions are often very complicated. Furthermore, these instances of wagering mechanisms do not admit a dominant strategy; rather, players make reports conditionally on their beliefs on other players' reports. At best, when the range of possible reports is limited to outcomes or events, the mechanism may allow a partial specification of individual subjective probabilities such as: the (subjective) probability is greater than some threshold. However this is an incomplete specification and less natural than simply stating a belief.

2.3 Distribution Properties

Distribution properties, introduced in Lambert et al. [11], are a convenient way to model information on probability distributions. A *distribution property* $\Gamma(P)$ is defined as a function that assigns a real value to any probability distribution P in a given convex domain. By assumption, the domain must contain the true distribution of the random experiment under consideration. For example, in the case of a continuous outcome, the domain of the median would be the set of continuous densities with full support, the domain of the expectation the set of distributions with finite first moment.

Common distribution properties include the probability of an event, the expectation, the variance, medians/quantiles, moments, indicators of the symmetry of the distribution (skewness), and dispersion (kurtosis). For instance, the property Γ corresponding to the probability of an event A is $\Gamma(P) = P(A)$, that of the median of a random variable X is $\Gamma(P) = \sup_m\{m : P(X < m) < 1/2\}$. As in Lambert et al. [11], we say that a report r (respectively, a probability P)

[1] The amount wagered may depend on the player's belief/knowledge. For example, the more confident the player is about some outcome, the more she may want to wager.

[2] This notational change is useful to define wagering mechanisms with a varying number of players.

is admissible when r (respectively $\Gamma(P)$) falls in the interior of Γ's range, which is an open interval for continuous properties. We are interested in wagering mechanisms wherein players report sets or vectors of property values. For example, a single probability of a given event, the full distribution for a finite set of outcomes, or a pair *(expectation, variance)*.

Similar to scoring rules [16], Lambert et al. [11] define reward functions that truthfully elicit single properties or sets of properties from a risk neutral agent as follows.

Definition 2. A score function for a vector of distribution properties $\Gamma = (\Gamma_1, \ldots, \Gamma_k)$ is a real-valued function $s(r, \omega)$, with $r = (r_1, \ldots, r_k)$ and r_i the report for property Γ_i, and ω the outcome. It is called *strictly proper* when

$$E_P[s(r, \omega)] < E_P[s(\Gamma(P), \omega)] \tag{1}$$

for all admissible probability P and vector $r = (r_1, \ldots, r_k) \neq (\Gamma_1(P), \ldots, \Gamma_k(P))$.

Here and throughout the remainder of the paper, $E_P[X]$ denotes the expectation of X when the outcome ω is distributed according to the distribution P. For convenience, we identify single properties as vectors of properties with a single element ($k = 1$).

3. WEIGHTED-SCORE MECHANISMS

In this section we present a specific subclass of wagering mechanisms called *weighted-score mechanisms*. These mechanisms are weighted mixtures of strictly proper score functions. Rewards are determined by (1) the relative performance of the players, as in the scoring rules of Kilgour and Gerchak [10], and (2) the amounts wagered, as in parimutuel betting markets.

Definition 3. A *weighted-score mechanism* is a wagering mechanism $(\Omega, \mathcal{R}, \Pi)$ associated with a vector of properties $\Gamma = (\Gamma_1, \ldots, \Gamma_k)$. Here Ω is the space of possible outcomes, $\mathcal{R} = (a_1, b_1) \times \cdots \times (a_k, b_k)$ where (a_i, b_i) is an interval of possible values for Γ_i, and Π is the vector of payout functions with the payout to bettor i defined as

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

where s is a strictly proper score function for Γ taking values in the interval $[0, 1]$.

We need to choose s so as verify Equation (1). Fortunately, there already exists a wealth of such functions for common properties [16]. Savage [14] gives a simple characterization of strictly proper score functions for probabilities and expectations of random variables, and examples of strictly proper scores for quantiles appear in Gneiting and Raftery [6]. For example, to elicit the probability of an event A, one may use the quadratic score $s(p, \omega) = 1 - (\mathbb{1}_{\omega \in A} - p)^2$. To elicit the median of a continuous variable in the interval $[\alpha, \beta]$, one can use $s(m, \omega) = 1 - |m - \omega|/(\beta - \alpha)$. Table 1 gives examples of weighted-score mechanisms to elicit the probability of a binary event, and the expectation and median of a continuous random variable.

Weighted-score mechanisms are interesting and valuable because they satisfy many desirable properties. These are discussed in detail in the next section. In Section 5, we will see that the core properties are verified by the slightly larger class of weighted-score mechanisms parameterized by the total money wagered, and that those are the *only* mechanisms satisfying these properties.

Note that when one elicits probabilities of a binary event (when $\Gamma(P) = P(A)$ for some event A), and when all wagers are equal, payouts are proportional to those given by the KG-scoring rules [10]. As we shall see in Section 5, the multiplicative factor, $(n-1)/n$ when there are n participants, is necessary to induce sybilproofness.

Weighted-score mechanisms can be used for *free* forecasting. They reveal truthful predictions of distribution properties from a group of experts, without requiring any outside subsidy. Betting is another natural field of application. Indeed, common betting mechanisms are limited in several ways. For example, parimutuel markets do not work properly when all bettors agree that one event is much more likely than its alternatives. These markets are also not applicable with continuous outcomes, such as date and time. Besides the limited range of bets allowed prevents the full exploitation of one's information. Weighted-score mechanisms overcome these issues.

4. MECHANISM PROPERTIES

4.1 Desirable Properties

We describe seven desirable properties for wagering mechanisms. Denote by $M = (\mathcal{R}, \Omega, \Pi)$ a wagering mechanism, and let Γ be a vector of properties $\Gamma = (\Gamma_1, \ldots, \Gamma_k)$.

The first three properties are adapted from the axioms proposed by Kilgour and Gerchak [10].

BUDGET-BALANCE M is *budget-balanced* if the market generates neither profit nor loss, i.e,

$$\sum_i \Pi_i(\mathbf{r}, \mathbf{m}, \omega) = \sum_i m_i,$$

for any vector of reports $\mathbf{r}$, any vector of money wagered $\mathbf{m}$, and any outcome ω.

ANONYMITY M is *anonymous* if the payouts do not depend on the identity of the players. For any permutation σ of $\mathbb{N}$, any player i, and any outcome ω,

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

TRUTHFULNESS M is *truthful* for Γ when players maximize their expected payout when reporting true property values. For any player i, any admissible probability P, any set of reports of others $\mathbf{r}_{-i}$ and any vector of wagers $\mathbf{m}$,

$$E_P[\Pi_i((\mathbf{r}_{-i}, \Gamma(P)), \mathbf{m}, \omega)] > E_P[\Pi_i((\mathbf{r}_{-i}, r_i), \mathbf{m}, \omega)]$$

is satisfied for all $r_i \neq \Gamma(P)$.

When a player's expected payout depends both on her own report and on reports of others, it is natural to interpret the player's payment as her performance relative to other players. The relative performance of a player should increase either when the player's absolute performance increases, or

Table 1: Some examples of weighted-score mechanisms.

Space of outcomes:	Information elicited:	Payout for player i:
Any space Ω	Probability of an event $A \subset \Omega$	$m_i + m_i \frac{\sum_j m_j (\mathbb{1}_{\omega \in A} - r_j)^2}{\sum_j m_j} - m_i (\mathbb{1}_{\omega \in A} - r_i)^2$
$\Omega = [\alpha, \beta]$	Expectation of the outcome	$m_i + m_i \frac{\sum_j m_j (\omega - r_j)^2}{(\alpha - \beta)^2 \sum_j m_j} - m_i \frac{(\omega - r_i)^2}{(\alpha - \beta)^2}$
$\Omega = [\alpha, \beta]$	Median of the outcome	$m_i + m_i \frac{\sum_j m_j \lvert\omega - r_j\rvert}{\lvert\alpha - \beta\rvert \sum_j m_j} - m_i \frac{\lvert\omega - r_i\rvert}{\lvert\alpha - \beta\rvert}$

when the absolute performance of another player decreases. This property is captured by *normality*.[3]

NORMALITY A mechanism M is *normal* if, for any admissible probability distribution P, if any player i changes her report, the changes of expected payouts Δ_j, with respect to P, of any other player j is null or of the opposite sign of the changes of expected payouts Δ_i of player i.

In electronic platforms, it is important to prevent players from manipulating identities. To reflect this requirement, we complement truthfulness by *sybilproofness*, which ensures that payouts remain unchanged as a subset of players with the same reports manipulate user accounts, either by merging accounts or creating fake identities, or by doing transfers of wagers between them.

SYBILPROOFNESS M is *sybilproof* if for any subset of players $\mathcal{S} \subset \mathbb{N}$, for any reports $\mathbf{r}$ with $r_i = r_j$ for $i, j \in \mathcal{S}$, for any vectors of wagers $\mathbf{m}, \mathbf{m}'$ such that $m_i = m_i'$ for $i \notin \mathcal{S}$ and $\sum_{i \in \mathcal{S}} m_i = \sum_{i \in \mathcal{S}} m_i'$, the following two conditions hold.

For all $i \notin \mathcal{S}$, for all ω,

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = \Pi_i(\mathbf{r}, \mathbf{m}', \omega) \ ,$$

and for all ω,

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

The next property deals with the participation constraint. *Individual rationality* ensures that players get a nonnegative expected profit with respect to their belief.

INDIVIDUAL RATIONALITY M is *individually rational* if for any player i, for any admissible probability distribution P, for any wager $m_i > 0$, there exists a report r_i^* such that for any wagers $\mathbf{m}_{-i}$, and reports $\mathbf{r}_{-i}$ of other players, the expected profit of i is nonnegative:

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

Milgrom and Stokey [12] proved the *no-trade theorem* which states that, under certain assumptions, agents should not trade with each other in a zero-sum market. The no-trade theorem requires that agents have the same prior, that they be Bayesian rational (performing Bayesian belief updates), and that their Bayesian rationality be common knowledge. Individual rationality as defined above relies on the assumption that players do not satisfy all three conditions required by the no-trade theorem, thus can serve as a participation constraint. Indeed, in practice, people do trade in zero-sum and even negative-sum markets, for example because they don't share the same priors, because they are over-confident or bounded rational, or because there are noisy traders in the market.

Our final property deals with the connection between wagers and profits. To increase participation incentives and larger wagers, it is preferable that a player making a positive expected profit under her own belief makes an even higher profit by increasing her wager. Similarly, a player who loses money on expectation should lose even more by raising her participation level. We call it *monotonicity*.

MONOTONICITY M is *monotonic* if for any player i, any admissible probability distribution P, any vector of reports $\mathbf{r}$, any vector of wagers $\mathbf{m}$, and any $M_i > m_i$, either

$$\begin{aligned} 0 &< E_P[\Pi_i(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega) - m_i] \\ &< E_P[\Pi_i(\mathbf{r}, (\mathbf{m}_{-i}, M_i), \omega) - M_i] \end{aligned}$$

or

$$\begin{aligned} 0 &> E_P[\Pi_i(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega) - m_i] \\ &> E_P[\Pi_i(\mathbf{r}, (\mathbf{m}_{-i}, M_i), \omega) - M_i] \ . \end{aligned}$$

4.2 Attributes of weighted-score mechanisms

We show that weighted-score mechanisms satisfy all the desired properties described above.

THEOREM 1. *All weighted-score mechanisms are:*

1. *Budget-balanced,*
2. *Anonymous,*
3. *Truthful,*
4. *Sybilproof,*
5. *Normal,*
6. *Individually rational,*
7. *and Monotonic.*

Besides, due to their linear form, weighted-score mechanisms are also group strategyproof, in the sense that they incentivize honest behavior not only at the individual level but also at the group level: a group of players maximizes its expected global payout only when each of its members reports true property values.

[3]Note that individual rationality, truthfulness and budget-balance do not necessarily imply normality for more than two players. For example, the mechanism of Section 6.1 verifies all three properties but is not normal.

PROOF. We prove each property separately as follows..

(1) Budget-Balance For any $\mathbf{r}$, $\mathbf{m}$, and ω,

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

(2) Anonymity Let σ be any permutation of $\mathbb{N}$. For any $\mathbf{r}$, $\mathbf{m}$, ω, and i,

$$\begin{aligned}&\Pi_{\sigma(i)}((r_{\sigma^{-1}(j)})_{j\in\mathbb{N}},(m_{\sigma^{-1}(j)})_{j\in\mathbb{N}},\omega) \\ &= m_{\sigma^{-1}(\sigma(i))}\Bigg(1+s(r_{\sigma^{-1}(\sigma(i))},\omega) \\ &\qquad -\frac{\sum_j s(r_{\sigma^{-1}(j)},\omega)m_{\sigma^{-1}(j)}}{\sum_j m_{\sigma^{-1}(j)}}\Bigg) \\ &= m_i\left(1+s(r_i,\omega)-\frac{\sum_j s(r_j,\omega)m_j}{\sum_j m_j}\right) \\ &= \Pi_i((r_j)_{j\in\mathbb{N}},(m_j)_{j\in\mathbb{N}},\omega) \ .\end{aligned}$$

(3) Truthfulness For any $\mathbf{r}$, $\mathbf{m}$, ω, i, and P,

$$\begin{aligned}E_P[\Pi_i(\mathbf{r},\mathbf{m},\omega)] &= m_i\Bigg(1+E_P[s(r_i,\omega)]\left(1-\frac{m_i}{\sum_j m_j}\right) \\ &\qquad -\frac{\sum_{j\neq i} E_P[s(r_j,\omega)]m_j}{\sum_j m_j}\Bigg).\end{aligned}$$

Since s is strictly proper for Γ, $E_P[s(r_i,\omega)]$ is maximized only at $r_i=\Gamma(P)$, so

$$E_P[\Pi_i((\mathbf{r}_{-i},r_i),\mathbf{m},\omega)] < E_P[\Pi_i((\mathbf{r}_{-i},\Gamma(P)),\mathbf{m},\omega)]$$

for all $r_i \neq \Gamma(P)$.

(4) Sybilproofness Let r be the common report of all $i \in \mathcal{S}$. For any $i \notin \mathcal{S}$,

$$\begin{aligned}&\Pi_i(\mathbf{r},\mathbf{m},\omega) \\ &= m_i\,(1+s(r_i,\omega) \\ &\qquad -\frac{\sum_{j\notin\mathcal{S}} s(r_j,\omega)m_j + s(r,\omega)\sum_{j\in\mathcal{S}} m_j}{\sum_{j\notin\mathcal{S}} m_j+\sum_{j\in\mathcal{S}} m_j} \\ &= m'_i\,(1+s(r_i,\omega) \\ &\qquad -\frac{\sum_{j\notin\mathcal{S}} s(r_j,\omega)m'_j + s(r,\omega)\sum_{j\in\mathcal{S}} m'_j}{\sum_{j\notin\mathcal{S}} m'_j+\sum_{j\in\mathcal{S}} m'_j} \\ &= \Pi_i(\mathbf{r},\mathbf{m}',\omega).\end{aligned}$$

Additionally,

$$\begin{aligned}&\sum_{i\in\mathcal{S}} \Pi_i(\mathbf{r},\mathbf{m},\omega) \\ &= \sum_{i\in\mathcal{S}} m_i\,(1+s(r,\omega) \\ &\qquad -\frac{\sum_{j\notin\mathcal{S}} s(r_j,\omega)m_j + s(r,\omega)\sum_{j\in\mathcal{S}} m_j}{\sum_{j\notin\mathcal{S}} m_j+\sum_{j\in\mathcal{S}} m_j} \\ &= \left(\sum_{i\in\mathcal{S}} m'_i\right)(1+s(r,\omega) \\ &\qquad -\frac{\sum_{j\notin\mathcal{S}} s(r_j,\omega)m'_j + s(r,\omega)\sum_{j\in\mathcal{S}} m'_j}{\sum_{j\notin\mathcal{S}} m'_j+\sum_{j\in\mathcal{S}} m'_j} \\ &= \sum_{i\in\mathcal{S}} \Pi_i(\mathbf{r},\mathbf{m}',\omega).\end{aligned}$$

(5) Normality Let $\tilde{r}_i$ and $\tilde{\mathbf{r}}=(\mathbf{r}_{-i},\tilde{r}_i)$ be defined in such a way that

$$E_P[\Pi_i(\tilde{\mathbf{r}},\mathbf{m},\omega)] > E_P[\Pi_i(\mathbf{r},\mathbf{m},\omega)].$$

Then

$$E_P[s(\tilde{r}_i,\omega)] > E_P[s(r_i,\omega)],$$

and for $j\neq i$,

$$\begin{aligned}&E_P[\Pi_j(\tilde{\mathbf{r}},\mathbf{m},\omega)] - E_P[\Pi_j(\mathbf{r},\mathbf{m},\omega)] \\ &= -\frac{m_i}{\sum_j m_j}(E_P[s(\tilde{r}_i,\omega)] - E_P[s(r_i,\omega)]) < 0.\end{aligned}$$

Similarly, if $E_P[\Pi_i(\tilde{\mathbf{r}},\mathbf{m},\omega)] < E_P[\Pi_i(\mathbf{r},\mathbf{m},\omega)]$ then $E_P[\Pi_j(\tilde{\mathbf{r}},\mathbf{m},\omega)] - E_P[\Pi_j(\mathbf{r},\mathbf{m},\omega)] > 0$. This proves normality.

(6) Individual rationality For fixed wagers $\mathbf{m}$ and a probability distribution P, $E_P[s(r_i,\omega)]$ is maximized when $r_i=\Gamma(P)$. Thus, $E_P[s(r_j,\omega)] \leq E_P[s(\Gamma(P),\omega)]$ for all j. We have

$$\frac{\sum_j E_P[s(r_j,\omega)]m_j}{\sum_j m_j} \leq E_P[s(\Gamma(P),\omega)].$$

Hence,

$$\begin{aligned}E_P[\Pi_i(\mathbf{r},\mathbf{m},\omega)-m_i] &= m_iE_P[s(r_i,\omega)] - m_i\frac{\sum_j E_P[s(r_j,\omega)]m_j}{\sum_j m_j} \\ &\geq m_i[E_Ps(r_i,\omega)] - m_iE_P[s(\Gamma(P),\omega)] \\ &= 0\end{aligned}$$

when player i reports $r_i=\Gamma(P)$.

(7) Monotonicity Let $\tilde{s}_i = E_P[s(r_i,\omega)]$. Then

$$E_P[\Pi_i(\mathbf{r},\mathbf{m},\omega)-m_i] = m_i\left(\tilde{s}_i - \frac{\sum_j \tilde{s}_jm_j}{\sum_j m_j}\right) = \frac{am_i}{m_i+b},$$

where $a=\tilde{s}_i\sum_{j\neq i} m_j - \sum_{j\neq i}\tilde{s}_jm_j$ and $b=\sum_{j\neq i} m_j$.

Because the value of $ax/(x+b)$ is positive and increases with x when $a>0$ and is negative and decreases with x when $a<0$, monotonicity holds. $\square$

5. UNIQUENESS OF WEIGHTED SCORE MECHANISMS

Theorem 1 shows that the family of weighted score mechanisms satisfies a number of useful properties. In this section, we will show that weighted score mechanisms are unique in this sense. More precisely, weighted score mechanisms, parameterized by the total money wagered in the common pool, are the *only* wagering mechanisms that are simultaneously budget-balanced, anonymous, truthful, normal, and sybilproof.

We start by characterizing the set of all truthful and normal wagering mechanisms, and progressively add the constraints of anonymity, budget-balance, and sybilproofness. This incidentally proves that the mechanisms introduced by Kilgour and Gerchak [10] without subsidy are the only *Competitive Prediction Schemes* to be truthful, normal, anonymous, and budget-balanced.

In the analysis that follows, we assume that Ω is finite, and consider elicitation of a single property Γ (for example, the probability of a binary event). We assume Γ is continuous and not locally constant,[4] and denote by (a, b) the interval of admissible reports. As in Lambert et al. [11], we define a distribution property to be elicitable when there exists a strictly proper score function for that property. We say that a function is smooth when it is twice continuously differentiable. In the sequel of this section, we consider mechanisms with smooth payouts, and when we refer to the term "wagering mechanism" we always mean "wagering mechanism with smooth payouts".

5.1 Characterizing truthful and normal mechanisms

The first characterization lemma shows that any wagering mechanism is truthful and normal when it is additively separable into strictly proper score functions.

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*

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

where for all i *and* j*, for any fixed value of* $\mathbf{m}$*,* $f_{i,j}$ *is a smooth strictly proper score function for* Γ*.*

The proof will make use of the following lemma, whose proof is omitted due to lack of space.

LEMMA 2. *If* $f : (a, b)^n \mapsto \mathbb{R}$ *is twice continuously differentiable, and if*

$$\frac{\partial^2 f(x_1, \ldots, x_n)}{\partial x_i \partial x_j} = 0$$

for all $\mathbf{x} \in \mathbb{R}^n$ *and all* $i \neq j$*, then there exists* $f_i : (a, b) \mapsto \mathbb{R}$ *such that*

$$f(x_1, \ldots, x_n) = \sum_{i=1}^{n} f_i(x_i) \ .$$

PROOF (LEMMA 1). The truthfulness and normality of any rule of this form follow from the linearity of expectation and $f_{i,j}$ being strictly proper scores for Γ. It remains to show that truthful and normal rule must be of this form.

Let n be the number of participating players. Since we assume the wagers $m_1, \cdots, m_n$ are fixed, we can denote the payout function for i as $\Pi_i(r_1, \ldots, r_n, \omega) = \Pi_i(\mathbf{r}, \omega)$.

Let P_i be the beliefs of bettor i. By truthfulness,

$$\Gamma(P_i) \in \arg\max E_{P_i}[\Pi_i(\mathbf{r}_{-i}, \cdot, \omega)] \tag{2}$$

for all $\mathbf{r}_{-i}$. By normality, for any $k \neq i$,

$$\Gamma(P_i) \in \arg\min E_{P_i}[\Pi_j(\mathbf{r}_{-i}, \cdot, \omega)] \ . \tag{3}$$

The first order condition in (2) and (3) gives that for all admissible beliefs P_i, for all $k \in \{1, \cdots, n\}$,

$$\left.\frac{\partial E_{P_i}[\Pi_k(\mathbf{r}, \omega)]}{\partial r_i}\right|_{r_i = \Gamma(P_i)} = 0 \ .$$

Applying the same argument to any $j \neq i$ with beliefs P_j gives us

$$\left.\frac{\partial E_{P_j}[\Pi_k(\mathbf{r}, \omega)]}{\partial r_j}\right|_{r_j = \Gamma(P_j)} = 0 \ .$$

[4]With respect to the topology considered in Lambert et al. [11].

We now differentiate the left side of the these equations by r_j and r_i respectively to get

$$\left.\frac{\partial^2 E_{P_i}[\Pi_k(\mathbf{r}, \omega)]}{\partial r_i \partial r_j}\right|_{r_i = \Gamma(P_i)} = \left.\frac{\partial^2 E_{P_j}[\Pi_k(\mathbf{r}, \omega)]}{\partial r_j \partial r_i}\right|_{r_j = \Gamma(P_j)} = 0 \ .$$

Now consider admissible P_i, P_j with different property values; in other words, $\Gamma(P_i) \neq \Gamma(P_j)$. Since Π_k is smooth,

$$\frac{\partial^2 \Pi_k}{\partial r_i \partial r_j} = \frac{\partial^2 \Pi_k}{\partial r_j \partial r_i} \ ,$$

so by linearity of expectation, for all admissible P_i and P_j,

$$E_{P_i}\left[\left.\frac{\partial^2 \Pi_k(\mathbf{r}, \omega)}{\partial r_i \partial r_j}\right|_{r_i = \Gamma(P_i), r_j = \Gamma(P_j)}\right] = 0 \ ,$$

and

$$E_{P_j}\left[\left.\frac{\partial^2 \Pi_k(\mathbf{r}, \omega)}{\partial r_i \partial r_j}\right|_{r_i = \Gamma(P_i), r_j = \Gamma(P_j)}\right] = 0 \ .$$

Let $r_i^* \neq r_j^*$ be admissible reports, and let X be the random variable defined by

$$X(\omega) = \left.\frac{\partial^2 \Pi_k(\mathbf{r}, \omega)}{\partial r_i \partial r_j}\right|_{r_i = r_i^*, r_j = r_j^*} \ .$$

The previous equations says that for all probabilities P_i, P_j such that $\Gamma(P_i) = r_i^*$ and $\Gamma(P_j) = r_j^*$, we have $E_{P_i}[X] = E_{P_j}[X] = 0$.

Now, Theorem 2 of Lambert et al. [11] shows that elicitable properties may be expressed as linear constraints. In particular, it shows that there exist random variables Y_i and Y_j such that

$$E_P[Y_i] = 0 \Leftrightarrow \Gamma(P) = r_i^*$$

and

$$E_P[Y_j] = 0 \Leftrightarrow \Gamma(P) = r_j^* \ .$$

Since $E_P[Y_i] = 0$ implies $\Gamma(P) = r_i^*$ which in turn implies that $E_P[X] = 0$, it must be the case that X is proportional to Y_i. By a similar argument, X must be proportional to Y_j. If X is not null, Y_i and Y_j are proportional, implying that $r_i^* = r_j^*$, which is a contradiction. Therefore it must be the case that $X(\omega) = 0$ for all ω, and for all admissible reports $r_i^* \neq r_j^*$,

$$\left.\frac{\partial^2 \Pi_k(r_1, \ldots, r_n, \omega)}{\partial r_i \partial r_j}\right|_{r_i = r_i^*, r_j = r_j^*} = 0 \ .$$

This remains true when $r_i^* = r_j^*$ by continuity.

By Lemma 2, this implies that there are functions $f_{k,i}$ such that

$$\Pi_k(\mathbf{r}, \omega) = f_{k,k}(r_k, \omega) - \sum_{i \neq k} f_{k,i}(r_i, \omega) \ .$$

By truthfulness, $f_{k,k}$ must be a strictly proper score function for Γ; by normality, for every $i \neq k$, $f_{k,i}$ must be a strictly proper score function for Γ. □

5.2 Adding Anonymity

With anonymity added, we provide a necessary condition for the special case of identical wagers.

LEMMA 3. *If a wagering mechanism is* truthful *for* Γ, normal*, and* anonymous*, then there exist smooth functions* f *and* g *such that if every player wagers the same amount* m*, then for all* i *the payout to* i *is*

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

where M *is the total amount wagered. Furthermore, for any fixed values of* m *and* M*,* f *and* g *must be strictly proper score functions for* Γ*.*

The proof, omitted due to space restrictions, is based on successive use of the anonymity property.

5.3 Adding budget-balance

We now add budget-balance, and easily show the following, still considering the special case of identical wagers.

LEMMA 4. *If a wagering mechanism is* truthful *for* Γ, normal*,* anonymous*, and* budget-balanced*, then there exists a smooth function* f *such that if every agent wagers the same amount* m*, then payouts are*

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

where M *is the total amount wagered. Furthermore, for any fixed values of* m *and* M*,* f *must be a strictly proper score function for* Γ*.*

The proof can be obtained directly by the application of the budget-balanced equality.

We can always write payouts as $\Pi_i = m_i(1 + \rho_i)$ for some function ρ_i, the Return On Investment (ROI). As we require nonnegative payouts, the function $f(\cdot, m, M, \omega)$ takes values in an interval of length at most m, so that the ROI of any player in such a mechanism is never higher than 1 when wagers are identical.

COROLLARY 1. *If a wagering mechanism is* budget-balanced*,* anonymous*,* truthful *and* normal*, a player's ROI cannot be above 100% when all wagers are identical.*

We will see that, adding sybilproofness, this corollary generalizes to the case of different wagers.

Interestingly, when applied to the setting of Kilgour and Gerchak, our result permits to prove uniqueness of their scoring rules within the class of Competitive Prediction Schemes (CPS) [10]. They restrict themselves to the case of eliciting the probability of a binary event; here we consider the natural generalization of their setting for the elicitation of one or more distribution properties. A CPS is essentially a wagering mechanism without a wager. In a first step each player i among n players makes a report r_i corresponding to one or several distribution properties. In a second step, after the true outcome ω of the uncertain event becomes known, player i receives a payment $\Pi_i(r_1, \ldots, r_n, \omega)$ (which may be negative). The properties *budget-balance*, *anonymity*, *truthfulness*, *normality*, and *individual rationality* can be directly adapted to this setting. The (generalized) Kilgour and Gerchak scoring rules, defined by

$$\Pi_i(r_1, \ldots, r_n, \omega) = s(r_i, \omega) - \frac{1}{n-1} \sum_{j \neq i} s(r_j, \omega) \ ,$$

where s a strictly proper score function for a vector of properties Γ, can be shown to be budget-balanced, anonymous, truthful, normal, and individually rational (based on K&G [10] and Theorem 1). Noting that the constraint of nonnegative payouts plays no role in proving Lemma 1 and 3, Lemma 4 shows that these scoring rules are the *only* competitive prediction schemes to satisfy these core properties.

THEOREM 2. *Given* n *players, any Competitive Prediction Scheme is* budget-balanced*,* anonymous*,* truthful *for* Γ *and* normal *if and only if the payment of agent* i *when the true outcome is* ω *is given by*

$$s(r_i, \omega) - \frac{1}{n-1} \sum_{j \neq i} s(r_j, \omega) \ ,$$

where s *is strictly proper for* Γ*.*

5.4 Uniqueness of Weighted Score

We now turn to the main theorem.

THEOREM 3. *A wagering mechanism is* budget-balanced*,* anonymous*,* truthful *for* Γ*,* normal *and* sybilproof*, if and only if the payouts are given by*

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

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

This result allows to complement Corollary 1.

COROLLARY 2. *If a wagering mechanism is* budget-balanced*,* anonymous*,* truthful*,* normal*, and* sybilproof*, a player's ROI is never above 100%.*

PROOF (THEOREM 3). The proof of Theorem 1 may be applied directly to show that weighted score mechanisms satisfy these properties. Here we show the other direction. The proof proceeds in three steps below.

To start, suppose that there are n players wagering the same amount m, and let $M = nm$ be the total amount wagered. Since the wagers are identical, we know that the payout function takes the special form given in Lemma 4. Let $\tilde{f}(r, m, M, \omega)$ be the function f given by this lemma. Fix any possible report r_0 and let $f(r, m, M, \omega) = \tilde{f}(r, m, M, \omega) - \tilde{f}(r_0, m, M, \omega)$, then

$$\begin{aligned} & \Pi_i(\mathbf{r}, \mathbf{m}, \omega) \\ & = \quad m + f(r_i, m, M, \omega) - \frac{1}{n-1} \sum_{j \neq i} f(r_j, m, M, \omega), \quad (4) \end{aligned}$$

and $f(r_0, m, M, \omega) = 0$ for all m, M, ω.

In the first step, we start with this equation and show that it is possible to create a function s such that for any positive integers k and n,

$$s\left(r, \omega, \frac{n}{2^k}\right) = 2^k \frac{n}{n-1} f\left(r, \frac{1}{2^k}, \frac{n}{2^k}, \omega\right) \ . \quad (5)$$

In the second step, we show that for any number of players n, any set of reports $\mathbf{r}$, and any set of (not necessarily identical) wagers $\mathbf{m}$, we can write the payout to player i as

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m_i \left(1 + s(r_i, \omega, M) - \frac{\sum_j m_j s(r_j, \omega, M)}{M} \right) \quad (6)$$

where $M = \sum_i m_i$ and s is the function defined in Step 1.

Finally, in Step 3, we show that the function s can be written as a strictly proper score function taking values in $[0, 1]$, completing the proof. More details on each step follow.

STEP 1:

Let (a, b) be the interval of possible reports. We begin by showing how to create a function $s : (a, b) \times \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ that satisfies Equation (5) for any positive integers k and n. To do so, we consider z groups of 2^ℓ players, for some $\ell > 0$. Every player in each of the groups wagers $1/(2^{k+\ell})$. In the first group, each player reports r, while in the other groups, each player reports r_0. By summing over the value of Equation (4) for each member of the first group, we see that the aggregate payout for the first group is

$$\frac{1}{2^k} + 2^\ell \left(1 - \frac{2^\ell - 1}{z \cdot 2^\ell - 1}\right) f\left(r, \frac{1}{2^{k+\ell}}, \frac{z}{2^k}, \omega\right) \ . \tag{7}$$

Now consider an alternate situation in which a single player wagers $1/2^k$ and reports r against $n-1$ other players, each wagering $1/2^k$ and reporting r_0. By Equation (4), the payout of the first player in this scenario is

$$\frac{1}{2^k} + f\left(r, \frac{1}{2^k}, \frac{n}{2^k}, \omega\right) \ . \tag{8}$$

By sybilproofness, Equation (7) must equal Equation (8) when the number of groups z is n. Hence by simple algebra,

$$\begin{aligned} 2^k \left(\frac{n - 2^{-\ell}}{n-1}\right) f\left(r, \frac{1}{2^k}, \frac{n}{2^k}, \omega\right) & \\ = \ 2^{k+\ell} f\left(r, \frac{1}{2^{k+\ell}}, \frac{n}{2^k}, \omega\right) \ . & \end{aligned} \tag{9}$$

Now, since

$$\lim_{\ell \to +\infty} 2^k \left(\frac{n - 2^{-\ell}}{n-1}\right) = 2^k \left(\frac{n}{n-1}\right) \ ,$$

it must be the case that the limit of the right-hand side of Equation (9) exists. Thus there exists a function $s : (a, b) \times \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ such that

$$s\left(r, \omega, \frac{n}{2^k}\right) = \lim_{\ell \to +\infty} 2^{k+\ell} f\left(r, 2^{-(k+\ell)}, \frac{n}{2^k}, \omega\right) \ .$$

Plugging this into Equation (9) yields Equation (5).

STEP 2:

Using the previous result, we now show that for any number of players n, for any reports and (not necessarily identical) wagers, the payout to each player i can be expressed as in Equation (6). To do so we use a continuity argument applied to an approximation of the payouts. In particular, we assume that the vector of wagers belongs to the set

$$\mathcal{M} = \{(a_1 2^{-k}, \ldots, a_n 2^{-k}) / \mathbf{a} \in \mathbb{N}^n, k \in \mathbb{N}\} \ .$$

While this assumption does restrict the set of possible wagers, we note that because $\mathcal{M}$ is dense in $\mathbb{R}^n_+$, it is the case that for *any* positive vector of wagers $\mathbf{m}$, there exist elements of $\mathcal{M}$ arbitrarily close to $\mathbf{m}$. Thus any vector of wagers can be approximated arbitrarily well using a vector in $\mathcal{M}$.

Now, given a vector of wagers $\mathbf{m} \in \mathcal{M}$, there clearly exist a vector of integers $\mathbf{a}$ and a value $k > 0$ such that $m_i = a_i/2^k$ for all i. Let $N = \sum_i a_i$. To obtain the vector of payouts $\Pi(\mathbf{r}, \mathbf{m}, \omega)$, we compare the present scenario with another in which there are N players divided into n groups. Each group i contains a_i players, each wagering $1/2^k$ on report r_i. By Equation (4) and (5) the aggregate payout of group i in this alternate scenario is

$$\begin{aligned} & \frac{a_i}{2^k} + a_i \Bigg(\left(1 - \frac{a_i - 1}{N-1}\right) f\left(r_i, \frac{1}{2^k}, \frac{N}{2^k}, \omega\right) \\ & \qquad - \frac{1}{N-1} \sum_{j \neq i} a_j f\left(r_j, \frac{1}{2^k}, \frac{N}{2^k}, \omega\right) \Bigg) \\ = \ & \frac{a_i}{2^k} + a_i \Bigg(\frac{N - a_i}{2^k N} s(r_i, \omega, M) \\ & \qquad - \sum_{j \neq i} \frac{a_j}{2^k N} s(r_j, \omega, M) \Bigg) \\ = \ & \frac{a_i}{2^k} \left(1 + s(r_i, \omega, M) - \sum_j \frac{a_j}{N} s(r_j, \omega, M)\right) \\ = \ & m_i \left(1 + s(r_i, \omega, M) - \frac{1}{M} \sum_j m_j s(r_j, \omega, M)\right) . \end{aligned}$$

By sybilproofness, this quantity is precisely $\Pi_i(\mathbf{r}, \mathbf{m}, \omega)$, the payout of player i in the initial scenario with n players. Thus since the set $\mathcal{M}$ is dense in $\mathbb{R}^n_+$, by continuity of the payout functions, for all $\mathbf{m} \in \mathbb{R}^n_+$, Equation (6) holds.

STEP 3:

It remains to show that the payout can always be written as in Equation (6) using a function s that is a score function taking values in $[0, 1]$. First, notice that $s(\cdot, \omega, M)$ must be bounded, since payouts are always nonnegative. Let $L(\omega, M) = \inf s(\cdot, \omega, M)$, and let $\tilde{s}(r, \omega, M) = s(r, \omega, M) - L(\omega, M)$. Note that we can write

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

Consider the scenario in which there are only two players, one wagering ϵ and reporting r, and the other wagering $M - \epsilon$ and reporting r'. The payout of the first player is

$$\epsilon \left[1 + \tilde{s}(r, \omega, M) - \frac{\epsilon \tilde{s}(r, \omega, M) + (M - \epsilon) \tilde{s}(r', \omega, M)}{M}\right] \geq 0 \ .$$

Since $\inf \tilde{s}(\cdot, \omega, M) = 0$, for all r',

$$1 - \frac{(M - \epsilon) \tilde{s}(r', \omega, M)}{M} \geq 0 \ ,$$

and, by taking the limit as $\epsilon \to 0$, for all r',

$$\tilde{s}(r', \omega, M) \leq 1 \ .$$

Therefore $\tilde{s}$ takes values in $[0, 1]$. The requirement of truthfulness implies that, for all $M > 0$, $\tilde{s}(\cdot, \cdot, M)$ is strictly proper for Γ, which concludes the proof. $\square$

6. EXTENSIONS

We now present several extensions of the weighted-score mechanisms, each of which achieves properties that weighted-score does not at the expense of other properties.

6.1 Adaptive weighted-score

One major difficulty encountered in designing score functions is that of incentive calibration: scores should vary the most in regions that are more likely to contain the true property value [14, 16]. As an example, consider designing a score function to collect reports about expected points in sports games. Without precise knowledge of the teams, the game, and its statistics, one must consider a large interval of possible point values. Yet informed forecasters are likely to report points within a small window. This mismatch induces small reward differences amongst forecasters. It also reduces incentives for agents to reveal their own belief with precision. However given expert advice limiting the plausible property values in advance, a score function with a large reward range for forecasters can be created.

The problem of properly adjusting score functions is particularly important for the weighted-score mechanism. Indeed, if score differences are small, the amplitude of money transfers between participants is likely to be considerably lower than initial wagers, thereby reducing incentives to participate.

We propose of a variant of weighted-score mechanism with self-adjusting score functions. To do this, consider a decomposition of the initial game into multiple, smaller games where only half of the players participate with a small fraction of their wagers. For each of these games, the payouts are computed according to a weighted-score payout function whose score functions are parameterized by sets of property values. This set, sent from the other half of the players, provides very accurate information about the regions where most reports are made. In any of the smaller games, players who participate do not influence the shape of the payout function, so that truthfulness remains true, although sybilproofness is lost.

We now give the formal definition. Given a set of n players $\mathcal{N}$, we consider the set of groups of players of size $\lceil n/2 \rceil$, $\mathcal{S} = \{A \subseteq \mathcal{N} / |A| = \lceil n/2 \rceil\}$, and $\mathcal{S}_i$ the set of sets of $\mathcal{S}$ containing i. For any set of property values R (provided, for examples, by experts as likely property values), let s^R be a strictly proper score function for Γ. Let $m'_i = m_i / |\mathcal{S}_i|$ be the wager of player i in each "small game". We call *adaptive* weighted-score mechanism a wagering mechanism whose payouts are

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = m_i + \sum_{S \in \mathcal{S}_i} m'_i \left[s^{\{r_\ell / \ell \in \mathcal{N} \setminus S\}}(r_i, \omega) - \frac{\sum_{j \in S} m'_j s^{\{r_\ell / \ell \in \mathcal{N} \setminus S\}}(r_j, \omega)}{\sum_{j \in S} m'_j} \right]$$

Compared to the original weighted-score mechanism, this variant loses normality and sybilproofness, but maintains the other properties.[5]

THEOREM 4. *Adaptive weighted-score mechanisms are* budget-balanced*,* anonymous*,* truthful*,* individually rational *and* monotonic*.*

6.2 Higher stakes

The notion of "betting m dollars" is commonly interpreted as placing a bet such that, in the worst case, m dollars will be lost. In the weighted-score mechanisms discussed here, players may not be able to lose their entire wager, no matter what the other players do. Indeed, when a player reports a value that does not minimize the score function, no matter what the outcome, they are guaranteed to recoup part of their wager. An important case is that of eliciting the probability of a binary event: Theorem 3 demonstrates that these "low stakes" are true for any mechanism having the core properties; any player announcing a probability other than 0 or 1 is certain to recoup part of her wager. Randomization can address this drawback.

Let $\Gamma = (\Gamma_1, \ldots, \Gamma_k)$ be a vector of distribution properties, s be strictly proper for Γ, and $\tilde{s}$ be proper (but not necessarily strictly proper), that is,

$$E_P[\tilde{s}(r, \omega)] \leq E_P[\tilde{s}(\Gamma(P), \omega)]$$

for all admissible reports r and probabilities P.

Consider computing the payout function as follows. First, flip a coin. If the coin is heads, then for all i,

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

Otherwise, for all i,

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

Randomization is equivalent to inserting a factor in the outcome space, using $\Omega' = \Omega \times \{H, T\}$, where $\{H, T\}$ is the outcome space for the coin flip. As long as the probabilities of heads and tails are strictly positive, the randomized mechanism conserves all of the properties of weighted-score wagering.[5]

THEOREM 5. *The randomized weighted-score mechanism is* budget-balanced*,* anonymous*,* truthful*,* individually rational*,* sybilproof*,* normal *and* monotonic*.*

If the alternative score only takes extreme values, a player can always lose her wager.[6] In the case of eliciting the probability of an event, one may choose

$$\tilde{s}(p, \omega) = \begin{cases} 1 & \text{if } |p - \omega| \leq \frac{1}{2} \ , \\ 0 & \text{otherwise.} \end{cases}$$

6.3 Higher ROI

By Corollaries 1 and 2, no wagering mechanism satisfying the core properties can result in returns on investment higher that 100%. This contrasts with most betting markets, wherein players may win several times their wager by betting on a very unlikely outcome. To solve this, one may use parimutuel-score mechanisms, where payouts are computed according to

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) = \frac{m_i s(r_i, \omega)}{\sum_j m_j s(r_j, \omega)} \sum_j m_j$$

where s is a strictly proper score function for Γ.

Parimutuel-score mechanisms are very similar to certain mechanisms introduced by Johnstone [8] and may be used in the same context. Like horse-race type parimutuel betting markets, parimutuel-score schemes have unbounded return on investment. They conserve several desirable properties.[5]

[5]Theorem 4, 5 and 6 may be proved in a similar fashion as Theorem 1. We omit the proofs due to space restrictions.

[6]At the limit as the total wager of other players grows.

Theorem 6. *Parimutuel-score mechanisms are* budget-balanced*,* anonymous*,* sybilproof *and* monotonic*.*

However, they are *not truthful*, *not individually rational* and *not normal*. They may be considered as approximately truthful in many practical situations: indeed, asymptotically as the number of players grows, a player tends to lose influence over the denominator of the payout function. The payout then becomes proportional to the strictly proper score function s, which ensures truthfulness.

7. CONCLUSION

We have investigated wagering mechanisms for revealing individual predictions from a group of agents. Agents are called to report on some information about some random experiment. Along with their report, they place a wager in a common pot. Upon realization of the experiment, agents receive a payment that depends on the true outcome and their own report. Payments are composed of a share of the common pot and possibly bonuses or taxes/fees. These mechanisms include many instances of common betting markets and call markets.

We have identified a particular subclass of such mechanisms called weighted-score mechanisms. Those novel elicitation schemes provide free individual forecasts of distribution properties, such as probabilities of binary events, or expectations and quantiles of random variables. They satisfy a number of desirable properties, including budget-balance, anonymity, truthfulness, normality, sybilproofness, individual rationality, and monotonicity.

Furthermore, we have showed that weighted-score mechanisms, parameterized by the total money wagered, are the only mechanisms that are budget-balanced, anonymous, truthful, normal, and sybilproof. In addition, we have proved that the Kilgour and Gerchak's scoring rules are the only forecasting methods to satisfy the core properties (for the setting considered in [10]).

We have explored variants of weighted-score mechanisms that conserve many desirable properties, and improve the core mechanism in several ways, by offering adaptive rewards, higher stakes, and unbounded returns on investment.

Our theoretical investigation leaves several open paths for future research. An important question, not addressed in the present work, is that of empirical studies. How do our mechanisms perform in a realistic environment? How do people behave when faced with a weighted-score mechanism or one of its variants? Also, our paper concentrated on one-shot mechanisms. A natural important step is to introduce dynamism: can we develop similar mechanisms – and similar characterizations – when adding a time dimension? In particular, is it possible to conserve truthfulness (and so avoid bluffing strategies), and incentivize agents to reveal their prediction early? This would permit to elicit at all times individual beliefs of agents, and watch their aggregation.

Acknowledgments

The authors would like to thank Esteban Arcaute, Alina Beygelzimer, Preston McAfee, Bethany Soule and Yevgeniy Vorobeychik for helpful discussions and suggestions, and the anonymous reviewers for their useful comments.

8. REFERENCES

[1] M. Ali. Probability and Utility Estimates for Racetrack Bettors. *The Journal of Political Economy*, 85(4):803–815, 1977.

[2] Y. Chen, L. Fortnow, E. Nikolova, and D. Pennock. Betting on permutations. *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 326–335, 2007.

[3] Y. Chen, D. Reeves, D. Pennock, R. Hanson, L. Fortnow, and R. Gonen. Bluffing and strategic reticence in prediction markets. In *WINE*, pages 70–81, 2007.

[4] N. Economides and J. Lange. A Parimutuel Market Microstructure for Contingent Claims Trading. *European Financial Management*, 11(1):25–49, 2005.

[5] L. Fortnow, J. Kilian, D. Pennock, and M. Wellman. Betting Boolean-style: a framework for trading in securities based on logical formulas. *Decision Support Systems*, 39(1):87–104, 2005.

[6] T. Gneiting and A. Raftery. Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

[7] R. Hanson. Combinatorial Information Market Design. *Information Systems Frontiers*, 5(1):107–119, 2003.

[8] D. Johnstone. The Parimutuel Kelly Probability Scoring Rule. *Decision Analysis*, 4(2):66, 2007.

[9] C.-H. Kehr, J. P. Krahnen, and E. Theissen. The anatomy of a call market. *Journal of Financial Intermediation*, 10(3-4):249–270, 2001.

[10] D. Kilgour and Y. Gerchak. Elicitation of Probabilities Using Competitive Scoring Rules. *Decision Analysis*, 1(2):108–113, 2004.

[11] N. Lambert, D. Pennock, and Y. Shoham. Eliciting Properties of Probability Distributions. *Proceedings of the 9th ACM Conference on Electronic Commerce*, 2008.

[12] P. Milgrom and N. Stokey. Information, trade, and common knowledge. *Journal of Economic Theory*, 1982.

[13] M. Peters, A. Man-Cho So, and Y. Ye. Pari-mutuel markets: Mechanisms and performance. In *WINE*, pages 82–95, 2007.

[14] L. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.

[15] R. H. Thaler and W. T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.

[16] R. Winkler, J. Muñoz, J. Cervera, J. Bernardo, G. Blattenberger, J. Kadane, D. Lindley, A. Murphy, R. Oliver, and D. Ríos-Insua. Scoring rules and the evaluation of probabilities. *TEST*, 5(1):1–60, 1996.

[17] J. Wolfers and E. Zitzewitz. Prediction Markets. *The Journal of Economic Perspectives*, 18(2):107–126, 2004.

[18] J. Wolfers and E. Zitzewitz. Prediction Markets in Theory and Practice. *Working paper*, 2006.