

Randomized Wagering Mechanisms*

Yiling Chen,¹ Yang Liu,² Juntao Wang¹

¹Harvard University, ²University of California, Santa Cruz
yiling@seas.harvard.edu, yangliu@ucsc.edu, juntaowang@g.harvard.edu

Abstract

Wagering mechanisms are one-shot betting mechanisms that elicit agents' predictions of an event. For deterministic wagering mechanisms, an existing impossibility result has shown incompatibility of some desirable theoretical properties. In particular, Pareto optimality (no profitable side bet before allocation) can not be achieved together with weak incentive compatibility, weak budget balance and individual rationality. In this paper, we expand the design space of wagering mechanisms to allow randomization and ask whether there are randomized wagering mechanisms that can achieve all previously considered desirable properties, including Pareto optimality. We answer this question positively with two classes of randomized wagering mechanisms: i) one simple randomized lottery-type implementation of existing deterministic wagering mechanisms, and ii) another family of randomized wagering mechanisms, named *surrogate wagering* mechanisms, which are robust to noisy ground truth. Surrogate wagering mechanisms are inspired by an idea of learning with noisy labels (Natarajan et al. 2013) as well as a recent extension of this idea to the information elicitation without verification setting (Liu and Chen 2018). We show that a broad set of randomized wagering mechanisms satisfy all desirable theoretical properties.

1 Introduction

Wagering mechanisms (Lambert et al. 2008; 2015; Chen et al. 2014; Freeman, Pennock, and Vaughan 2017; Freeman and Pennock 2018) are one-shot betting mechanisms that allow a principal to elicit participating agents' beliefs about an event of interest without paying out of pocket or incurring a risk. Compared with prediction-market type of dynamic elicitation mechanisms, one-shot wagering is possibly preferred due to its simplicity. It is particularly designed

for agents with immutable beliefs who “agree to disagree” and who do not update their beliefs. In a wagering mechanism, each agent submits a prediction for the event and specifies a wager, which is the maximum amount of money that the agent is willing to lose. Then after the event outcome is revealed, the total wagered money will be redistributed among the participants. Researchers have developed wagering mechanisms with various theoretical properties. In particular, Lambert et al. (2008; 2015) proposed a class of weighted score wagering mechanisms (WSWM) that satisfy a set of desirable properties, including budget balance, individual rationality, incentive compatibility, sybilproofness, among others.¹ Chen et al. (2014) later proposed a no-arbitrage wagering mechanism (NAWM) that removes opportunities for participating agents to risklessly profit.

However, in both WSWM and NAWM, it has been observed that a participant only loses a very small fraction of his total wager even in the worst case. This seems to be undesirable in practice as it is against the spirit of betting and a wager effectively loses its meaning as a budget. Freeman, Pennock, and Vaughan (2017) first formalized this observation by indicating that these mechanisms are not Pareto optimal, where Pareto optimality requires that there is no profitable side bet among participants before the allocation of a wagering mechanism is realized. They also proved an impossibility result: Pareto optimality cannot be satisfied together with individual rationality, weak budget balance and weak incentive compatibility. A double clinching auction (DCA) wagering mechanism (Freeman, Pennock, and Vaughan 2017) was hence proposed to improve Pareto efficiency. The parimutuel consensus mechanism (PCM) has been shown to satisfy Pareto optimality (Freeman and Pennock 2018), but violates incentive compatibility.

This paper is another quest of wagering mechanisms with better theoretical properties. We expand the design space of wagering mechanisms to allow randomization on agent payoffs and ask whether we can achieve all aforementioned desirable properties, including Pareto optimality. We give a positive answer to this question: Our randomized wagering mechanisms are the first ones to achieve Pareto optimality along with other properties.

We first show that a simple randomized lottery-type im-

*This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2017-17061500006. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Definitions of some properties can be found in Section 3.

plementation of existing wagering mechanisms (Lottery Wagering Mechanisms (LWM)) satisfy all desirable properties. In LWM, instead of receiving re-allocated money from a deterministic wagering mechanism, each agent receives a number of lottery tickets proportional to his payoff in the deterministic wagering mechanism. Then, the agent with the winning lottery wins the total wager (collected from all participants).

We then design another family of randomized wagering mechanisms, the Surrogate Wagering Mechanisms (SWM), by bringing insights from learning with noisy data (Natarajan et al. 2013; Scott 2015) to wagering mechanism design. A SWM first generates a “surrogate outcome” for each agent according to the true event outcome. An agent’s reported prediction is then evaluated using his surrogate but biased outcome together with a bias removal procedure such that in expectation the agent receives a score as if his prediction is evaluated against the true event outcome. Despite being randomized, SWM preserve the incentive properties of a deterministic wagering mechanism. We show that certain SWM satisfy all desirable properties of a wagering mechanism. Notably, SWM are robust to situations where only a noisy copy of the event outcome is available - this property is due to the fact that we borrow the machinery from the literature of learning with noisy data in designing SWM. We believe that this is another unique contribution to the literature of wagering mechanism design.

The rest of this paper is organized as follows. We discuss related work in the rest of this section. Section 2 introduces some preliminaries. We define randomized wagering mechanisms as well as desirable theoretical properties for them in Section 3. Section 4 presents a family of lottery-based wagering mechanisms. A family of surrogate wagering mechanisms are introduced in Section 5. Extensive simulations are conducted in Section 6 to demonstrate the advantages of randomized wagering mechanisms. Section 7 concludes this paper. Missing definitions and proofs can be found in the full version of the paper (Chen, Liu, and Wang 2018).

Related work

The ability to elicit *information*, in particular predictions and forecasts about future events, is crucial for many application settings and has been studied extensively in the literature. Proper scoring rules have been designed (Brier 1950; Jose, Nau, and Winkler 2006; Matheson and Winkler 1976; Winkler 1969; Gneiting and Raftery 2007) for this purpose. Later, competitive scoring rule (Kilgour and Gerchak 2004) and a parimutuel Kelly probability scoring rule (Johnstone 2007) adapt proper scoring rules to group competitive betting. Both mechanisms are budget balanced so that the principal doesn’t need to pay any participant. These spur the further development of the previously discussed wagering mechanisms (Lambert et al. 2008; 2015; Chen et al. 2014; Freeman, Pennock, and Vaughan 2017; Freeman and Pennock 2018) and the examination of their theoretical properties.

The idea of using randomization in wagering mechanism design is not entirely new. It first appeared in (Lambert et al. 2008), but not thoroughly studied. They restricted random-

ization to randomly selecting scoring rules to increase the maximum amount of money an agent can lose in WSWM. Cummings, Pennock, and Vaughan (2016) proposed to apply differential privacy technology to randomize the payoff of wagering mechanisms to preserve the privacy of agents’ beliefs. Our specific ideas of adding randomness are inspired by recent works on forecasting competition (Witkowski et al. 2018), surrogate scoring rules (Liu and Chen 2018), and the literature on learning with noisy labels (Bylander 1994; Natarajan et al. 2013; Scott 2015).

2 Preliminaries

In this section, we explain the scenario where a wagering mechanism applies and formally introduce the deterministic wagering mechanisms. Consider a scenario where a principal is interested in eliciting subjective beliefs from a set of agents $\mathcal{N} = \{1, 2, \dots, N\}$ about a random variable (event) X , which takes a value (outcome) in set $\mathcal{X} = \{0, 1, \dots, M - 1\}$, $M \geq 2$. The belief of each agent i is private, denoted as a vector of occurrence probabilities of each outcome $\mathbf{p}_i = (p_i^j)_{j \in \mathcal{X}} \in \Delta^{M-1}$. Following the previous work on wagering mechanism, this paper continues to adopt an immutable belief model for agents. Unlike in a Bayesian model, agents with immutable beliefs do not update their beliefs. The immutable belief model and the Bayesian model are two extremes of agent modeling for information elicitation, with the reality lies in between and arguably closer to the immutable belief side as people do “agree to disagree.” Moreover, Lambert et al. (2015) showed that while WSWM was designed for agents with immutable beliefs, it continued to perform well for Bayesian agents who have some innate utility for trading.

The principal uses a *wagering mechanism* to elicit private beliefs of agents. In a wagering mechanism, each agent reports a probability vector $\hat{\mathbf{p}}_i \in \Delta^{M-1}$, capturing his belief, and wagers an amount of money $w_i \in \mathbb{R}_+$. Similar to Lambert et al. (2008), we assume that wagers are exogenously determined for each agent and are not a strategic consideration. We use $\hat{\mathbf{p}}$ and \mathbf{w} to denote the reports and the wagers of all agents respectively, and use $\hat{\mathbf{p}}_{-i}$ and \mathbf{w}_{-i} to denote the reports and wagers of all agents other than agent i . In addition, we use $W_{\mathcal{S}}$ to denote $\sum_{i \in \mathcal{S}} w_i$ for any set of agents $\mathcal{S} \subseteq \mathcal{N}$. After an event outcome $x \in \mathcal{X}$ is realized, the wagering mechanism redistributes all the wagers collected from agents according to $\hat{\mathbf{p}}, \mathbf{w}, x$. The net-payoff of agent i is defined as the payoff or the money that agent i receives from the redistribution minus his wager. A wagering mechanism defines a net-payoff function $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$ for each agent i with wager constraint $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) \geq -w_i$ and constraint $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) = 0$ whenever $w_i = 0$. The two constraints ensure that no agent can lose more than his wager and no agent with zero wager can gain.

Strictly proper scoring rules and weighted score wagering mechanisms

Strictly proper scoring rules (Gneiting and Raftery 2007) are scoring functions proposed and developed to truthfully elicit beliefs from risk-neutral agents. They are building blocks

of many incentive compatible wagering mechanisms, such as **WSWM** and **NAWM**. A strictly proper scoring rule rewards a prediction $\hat{\mathbf{p}}_i$ by a score $s_x(\hat{\mathbf{p}}_i)$, according to the realization x of the random variable X . The scoring function $s_x(\cdot)$ is designed such that the expected payoff of truthful reporting is strictly larger than that of any other report, i.e., $\mathbb{E}_{X \sim \mathbf{p}_i} [s_X(\mathbf{p}_i)] > \mathbb{E}_{X \sim \mathbf{p}_i} [s_X(\hat{\mathbf{p}}_i)]$, $\forall \hat{\mathbf{p}}_i \neq \mathbf{p}_i$.

There is a rich family of strictly proper scoring functions, including Brier scores (for binary outcome event, $s_x(\hat{p}_i) = 1 - (\hat{p}_i - x)^2$, where \hat{p}_i is agent i 's report of $\mathbb{P}(X = 1)$), logarithmic and spherical scoring functions. Strictly proper scoring rules are closed under positive affine transformations.

WSWM (Lambert et al. 2008) rewards an agent according to his wager and the accuracy of his prediction relative to that of other agents' predictions. The net-payoff of agent i in **WSWM**, is formally defined as

$$\Pi_i^{\text{WS}}(\hat{\mathbf{p}}; \mathbf{w}; x) = \frac{w_i W_{\mathcal{N} \setminus \{i\}}}{W_{\mathcal{N}}} \left(s_x(\hat{\mathbf{p}}_i) - \sum_{j \in \mathcal{N} \setminus \{i\}} \frac{w_j}{W_{\mathcal{N} \setminus \{i\}}} s_x(\hat{\mathbf{p}}_j) \right), \quad (1)$$

where $s_x(\cdot)$ is any strictly proper scoring rule bounded within $[0, 1]$. **WSWM** strictly encourages truthful reporting of predictions, because the net-payoff of agent i is a strictly proper scoring rule of his prediction. Meanwhile, $\sum_{i \in \mathcal{N}} \Pi_i^{\text{WS}}$ is always zero by the form of the net-payoff formula, no matter what $s_x(\cdot)$ is. This means that the budget balance property of Eqn. (1) doesn't depend on the scoring rules. Our proposed surrogate wagering mechanisms use the same general form of the net-payoff function (but a different scoring rule) to guarantee the ex-post budget balance.

3 Randomized wagering mechanisms

We introduce randomized wagering mechanisms as extensions of deterministic wagering mechanisms. Similar to deterministic wagering mechanisms, the net-payoff of an agent in randomized wagering mechanisms depends on all agents' predictions $\hat{\mathbf{p}}$ and wagers \mathbf{w} , as well as the realized outcome x . But different from deterministic wagering mechanisms, the net-payoffs are now random variables. For notational simplicity, we now use $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$ to represent the random variable of agent i 's net-payoff in a randomized wagering mechanism. We use $\pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$ to represent the realization of $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$. We use Π_i and π_i as abbreviations when $\hat{\mathbf{p}}; \mathbf{w}; x$ are clear in the context. We denote the maximum/minimum possible value of a random variable X by \bar{X}/\underline{X} . We denote the joint distribution of $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$, $i \in \mathcal{N}$ by $\mathcal{D}(\hat{\mathbf{p}}; \mathbf{w}; x)$ and the marginal distribution of $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$ by $\mathcal{D}_i(\hat{\mathbf{p}}; \mathbf{w}; x)$.

Definition 1. Given a set \mathcal{N} of agents, reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} of agents and the event outcome x , a randomized wagering mechanism defines a joint distribution $\mathcal{D}(\hat{\mathbf{p}}; \mathbf{w}; x)$, and pays agent i by a net-payoff $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$, where $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$, $i \in \mathcal{N}$ are jointly drawn from $\mathcal{D}(\hat{\mathbf{p}}; \mathbf{w}; x)$. Moreover, $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) \geq -w_i$ and $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) = 0$ whenever $w_i = 0$.

A deterministic wagering mechanism is a special case of randomized wagering mechanisms when $\mathcal{D}_i(\hat{\mathbf{p}}; \mathbf{w}; x)$ is a point distribution for all agent $i \in \mathcal{N}$.

Desirable properties

In the literature, several desirable properties of wagering mechanisms have been proposed in the deterministic context. Lambert et al. (2008) introduced **(a)** individual rationality, **(b)** incentive compatibility, **(c)** budget balance, **(d)** sybilproofness, **(e)** anonymity, and **(f)** neutrality. Chen et al. (2014) introduced **(g)** no arbitrage. Freeman, Pennock, and Vaughan (2017) introduced **(h)** Pareto optimality. We extend these properties to the randomized context. These new properties reduce to the properties defined in the literature for the special case of deterministic wagering mechanisms.

(a) Individual rationality requires that each agent has nothing to lose in expectation by participating.

Definition 2. A randomized wagering mechanism is **individually rational (IR)** if $\forall i, \mathbf{p}_i, \mathbf{w}$, and $\hat{\mathbf{p}}_{-i}$, there exists $\hat{\mathbf{p}}_i$ such that

$$\mathbb{E}_{X \sim \mathbf{p}_i, \Pi_i \sim \mathcal{D}_i(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}; X)} [\Pi_i(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}; X)] \geq 0.$$

(b) Incentive compatibility requires that an agent's expected net-payoff is maximized when he reports honestly, regardless of other agents' reports and wagers.

Definition 3. A randomized wagering mechanism is **weakly incentive compatible (WIC)** if $\forall i, \mathbf{p}_i, \hat{\mathbf{p}}_i \neq \mathbf{p}_i, \hat{\mathbf{p}}_{-i}, \mathbf{w}$:

$$\mathbb{E}_{X \sim \mathbf{p}_i, \Pi_i \sim \mathcal{D}_i(\mathbf{p}_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}; X)} [\Pi_i(\mathbf{p}_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}; X)] \geq \mathbb{E}_{X \sim \mathbf{p}_i, \Pi_i \sim \mathcal{D}_i(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}; X)} [\Pi_i(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}; X)].$$

A randomized wagering mechanism is **strictly incentive compatible (SIC)** if the inequality is strict.

(c) Ex-post budget balance ensures that the principal does not need to subsidize the betting.

Definition 4. A randomized wagering mechanism is **weakly ex-post budget-balanced (WEBB)** if $\forall \hat{\mathbf{p}}, \mathbf{w}, x$: $\sum_{i \in \mathcal{N}} \pi_i(\hat{\mathbf{p}}, \mathbf{w}, x) \leq 0$ for any realization $(\pi_i)_{i \in \mathcal{N}}$ drawn from the joint distribution $\mathcal{D}(\hat{\mathbf{p}}, \mathbf{w}, x)$. A randomized wagering mechanism is **ex-post budget-balanced (EBB)** if the equality always holds.

(d) Sybilproofness requires that no agent can increase its expected net-payoff by creating fake identities and splitting his wager, regardless of other agents' reports and wagers.

Definition 5. Suppose agent i , instead of participating under one account with reported prediction $\hat{\mathbf{p}}_i$ and wager w_i , participates under $k > 1$ sybil accounts, with predictions and wagers $\{\hat{\mathbf{p}}_{i_l}, w_{i_l}\}_{l=1, \dots, k}$ such that $\hat{\mathbf{p}}_{i_l} = \hat{\mathbf{p}}_i, w_{i_l} \geq 0, \forall l = 1, \dots, k$ and $\sum_{l=1}^k w_{i_l} = w_i$. A randomized wagering mechanism is **sybilproof** if $\forall i, \hat{\mathbf{p}}, \mathbf{w}$, and x , and for all sybil reports $\hat{\mathbf{p}}_{i_1}, \dots, \hat{\mathbf{p}}_{i_k}$ and wagers w_{i_1}, \dots, w_{i_k} , we have

$$\mathbb{E}_{\Pi \sim \mathcal{D}(\hat{\mathbf{p}}; \mathbf{w}; x)} [\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)] \geq \mathbb{E}_{\Pi' \sim \mathcal{D}(\hat{\mathbf{p}}'; \mathbf{w}'; x)} \left[\sum_{l=1}^k \Pi_{i_l}(\hat{\mathbf{p}}_{i_l}; \mathbf{w}'_{i_l}; x) \right].$$

where $\hat{\mathbf{p}}$, \mathbf{w} and Π are the reports, wagers and net-payoffs when agent i participates under one account and $\hat{\mathbf{p}}'$, \mathbf{w}' and Π' are the reports, wagers and net-payoffs when agent i participates using k sybils.

(e) **Anonymity** requires that the distributions of random net-payoffs do not depend on agents' identities.

(f) **Neutrality** requires that the distributions of random net-payoffs do not depend on the labeling of the outcomes.

Formal definitions of properties (e) and (f) are presented in Section 4 of the full version (Chen, Liu, and Wang 2018).

(g) **No arbitrage** requires that no agent can risklessly make a profit.

Definition 6. A randomized wagering mechanism has **no arbitrage** if $\forall i, \hat{\mathbf{p}}, \mathbf{w} (\mathbf{w} > \mathbf{0}), \exists x$ such that $\underline{\Pi}_i(\hat{\mathbf{p}}, \mathbf{w}, x) < 0$.

(h) **Pareto optimality** in economics refers to an efficient situation where no trade can be made to improve an agent's payoff without harming any other agent's payoff. In an IR wagering mechanism, agents with different beliefs can always form a profitable (in expectation) wagering game if they all have a positive budget. Freeman, Pennock, and Vaughan (2017) defined Pareto optimality of a wagering mechanism as a property that agents with different beliefs will each lose all of his wager under at least one of the event outcomes. This "worst-case" outcome might be different for different agents. Thus, before the event outcome is realized, no agent can commit to secure part of his wager from the mechanism and no additional profitable wagering game can be made. We define Pareto optimality for randomized wagering mechanisms in a similar spirit: no agents with different beliefs can commit to secure part of their wagers before the event outcome is realized.

Definition 7. A randomized wagering mechanism is **Pareto optimal (PO)** if $\forall \hat{\mathbf{p}}, \mathbf{w}, \forall i, j \in \mathcal{N}$ with $\hat{\mathbf{p}}_i \neq \hat{\mathbf{p}}_j, \exists l \in \{i, j\}$ and x , such that $\underline{\Pi}_l(\hat{\mathbf{p}}, \mathbf{w}, x) = -w_l$.

Properties of existing wagering mechanisms

We summarize the properties of existing wagering mechanisms² in Table 1 in the full paper (Chen, Liu, and Wang 2018). No existing mechanism satisfies all properties (a)-(h). Moreover, Freeman, Pennock, and Vaughan (2017) showed an **impossibility result** that for deterministic wagering mechanisms, it is impossible to achieve properties IR, WIC, WEBB, and PO simultaneously. For existing randomized wagering mechanisms, the randomized WSWM in (Lambert et al. 2008) only satisfies PO in the limit of large population of participants, and the private WSWM (Cummings, Pennock, and Vaughan 2016) does not satisfy WEBB and PO.

4 Lottery wagering mechanisms

In this section we introduce a family of randomized wagering mechanisms, the *lottery wagering mechanisms (LWM)*,

²WSWM, NAWM, DCA, PCM, randomized WSWM (Lambert et al. 2008), private WSWM (Cummings, Pennock, and Vaughan 2016)

Mechanism 1 Lottery Wagering Mechanisms

- 1: Compute the payoff of each agent i under a DET: $\pi'_i \leftarrow w_i + \Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$.
 - 2: Each agent has winning probability $\frac{\pi'_j}{\sum_{i \in \mathcal{N}} \pi'_i}$. Draw a lottery winner $j^* \in \mathcal{N}$.
 - 3: Winner i^* is assigned a net-payoff $\sum_{i \in \mathcal{N} \setminus \{i^*\}} w_i$ and any agent $j \neq i^*$ has a net-payoff $-w_j$.
-

which extends arbitrary deterministic wagering mechanisms into randomized wagering mechanisms. We will show that LWM easily preserve (the randomized version of) the properties of the underlying deterministic wagering mechanisms, while achieving Pareto optimality, overcoming the impossibility result.

In lottery wagering mechanisms, each agent receives a number of lottery tickets in proportion to the *payoff* he gets under a deterministic wagering mechanism, and a winner is drawn from all the lottery tickets to win the entire pool of wagers. The mechanisms are designed in a way such that the expected payoff of each agent is equal to his payoff in the underlying deterministic wagering mechanisms and each agent has a positive probability to lose all his wager. Hence, no profitable side bet exists and the mechanisms are Pareto optimal. We formally present the lottery wagering mechanism that extends an arbitrary deterministic wagering mechanism DET in Mechanism 1. To distinguish the payoff from the net-payoff, we denote the payoff of agent i by π'_i .

Lottery wagering mechanisms are powerful in obtaining desirable theoretical properties. We show in Theorem 1 that the lottery wagering mechanism that extends WSWM, namely Lottery Weighted Score wagering mechanism (LWS), satisfies all properties (a)-(h).

Theorem 1. LWS satisfies all properties (a) - (h).

We notice that although LWS satisfies all desirable properties, it can be unsatisfying because (1) agents have high variance in payoff and (2) except the winning agent, all other agents lose money. To alleviate these issues, we can mix LWS with WSWM by assigning each of them a probability to be executed. The resulting mechanism still satisfies all the properties (a)-(h). The probabilistic mixture allows us to adjust the variance of the payoffs as well as agents' winning probabilities in the resulting mechanism.

5 Surrogate wagering mechanisms

In this section, we propose the *surrogate wagering mechanisms (SWM)*. We first introduce the generic SWM, then a variant of SWM that achieves the desirable theoretical properties and at the same time have moderate variance in payoffs and higher winning probabilities for accurate predictions. We then notice that randomization opens up the possibility of dealing with situations where only noisy ground truth is available. We discuss how to extend our results to this noisy setting.

Generic surrogate wagering mechanisms

A surrogate wagering mechanism consists of three main steps: (1) SWM generates a surrogate event outcome for each agent based on the true event outcome and a randomization device; (2) SWM evaluates each agent's prediction according to the surrogate event outcome using a designed scoring function such that the score is an unbiased estimate of the score derived by applying a strictly proper scoring rule to the ground truth outcome; (3) SWM applies WSWM to the scores based on the surrogate event outcome to determine the final net-payoff of each agent. Next, we explain these three steps in details. For clarity and simplicity of exposition, we consider only binary events, i.e., $\mathcal{X} = \{0, 1\}$, in this section.³

Step 1. Generating surrogate event outcomes A SWM generates a surrogate event outcome \tilde{X}_i for each agent $i \in \mathcal{N}$. Denote $\tilde{\mathbf{X}} = (\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_N)$. \tilde{X}_i 's are drawn independently conditional on X , and are specified by SWM. The conditionally marginal distribution $\mathbb{P}(\tilde{X}_i|X)$, $i \in \mathcal{N}$ can be expressed by two parameters, the error rates of the surrogate outcome: $e_1^i = \mathbb{P}(\tilde{X}_i = 0|X = 1)$ and $e_0^i = \mathbb{P}(\tilde{X}_i = 1|X = 0)$. The conditionally marginal distribution $\mathbb{P}(\tilde{X}_i|X)$ can be any distribution satisfying $\forall i \in \mathcal{N} : e_1^i + e_0^i \neq 1$.⁴ We use $\tilde{\mathbf{x}}$ and \tilde{x}_i to denote the realization of $\tilde{\mathbf{X}}$ and \tilde{X}_i respectively.

Step 2. Computing unbiased scores Given a strictly proper scoring rule $s_x(\cdot)$ within $[0,1]$, SWM computes the score of agent i as $\varphi \circ s_{\tilde{x}_i}(\hat{p}_i)$, where

$$\varphi \circ s_{\tilde{x}_i}(\hat{p}_i) = \frac{(1 - e_1^i - \tilde{x}_i)s_{\tilde{x}_i}(\hat{p}_i) - e_0^i s_{1-\tilde{x}_i}(\hat{p}_i)}{1 - e_0^i - e_1^i}. \quad (2)$$

\tilde{x}_i is the realized surrogate event outcome for agent i . Lemma 1 shows that φ is an unbiased operator on the score $s_{\tilde{x}_i}(p_i)$ in the sense that $\mathbb{E}_{\tilde{X}_i|x}[\varphi \circ s_{\tilde{X}_i}(\hat{p}_i)] = s_x(\hat{p}_i)$.

Lemma 1 (Lemma 3.4 of (Liu and Chen 2018)). $\forall x \in \{0, 1\}$, $\forall \hat{p}_i, e_0^i, e_1^i \in [0, 1]$ and $e_0^i + e_1^i \neq 1$, we have $\mathbb{E}_{\tilde{X}_i|x}[\varphi \circ s_{\tilde{X}_i}(\hat{p}_i)] = s_x(\hat{p}_i)$.

Lemma 1 implies that if $s_x(\hat{p}_i)$ is a strictly proper scoring rule, then $\varphi \circ s_{\tilde{x}_i}(\hat{p}_i)$ is also a strictly proper scoring rule.

Step 3. Computing net-payoffs In the final step, SWM computes the net-payoff of agent i using WSWM and the unbiased score of agent i , i.e., replacing score $s_x(\hat{p}_i)$ in Eqn. (1) by score $\varphi \circ s_{\tilde{x}_i}(\hat{p}_i)$. Formally, we have

$$\begin{aligned} \Pi_i^{\text{SWM}}(\hat{\mathbf{p}}, \mathbf{w}, x) &= \frac{w_i W_{\mathcal{N} \setminus \{i\}}}{W_{\mathcal{N}}} \left(\varphi \circ s_{\tilde{x}_i}(\hat{p}_i) \right. \\ &\quad \left. - \sum_{j \in \mathcal{N} \setminus \{i\}} \frac{w_j}{W_{\mathcal{N} \setminus \{i\}}} \varphi \circ s_{\tilde{x}_j}(\hat{p}_j) \right), \quad (3) \end{aligned}$$

³Extension to multi-outcome events can be found in Section 7.2 in the full version (Chen, Liu, and Wang 2018).

⁴If $e_0 + e_1 = 1$, \tilde{X}_i turns out to be independent with X , and thus provides no information about X . We thus exclude $e_1^i + e_0^i = 1$.

Mechanism 2 Surrogate Wagering Mechanisms

- 1: Collect the predictions $\hat{\mathbf{p}}$ and wagers \mathbf{w} .
- 2: Select error rate $e_0^i, e_1^i \in [0, 1]$ and $e_0^i + e_1^i \neq 1, \forall i$.
- 3: Generate surrogate outcome $\tilde{X}_i, \forall i$ such that $\mathbb{P}(\tilde{X}_i = 1|X = 0) = e_0^i, \mathbb{P}(\tilde{X}_i = 0|X = 1) = e_1^i$.
- 4: Score each agent $i \in \mathcal{N}$ according to Eqn. (2).
- 5: Pay each agent $i \in \mathcal{N}$ a net-payoff using Eqn. (3).

where x and $\tilde{x}_i, i \in \mathcal{N}$ are the event outcome and the surrogate event outcome for each agent i respectively.

We formally present SWM in Mechanism 2. According to Lemma 1 (applying to each score terms), we have $\forall i, x, \hat{\mathbf{p}}, \mathbf{w} : \mathbb{E}_{\Pi_i^{\text{SWM}} \sim \mathcal{D}(\hat{\mathbf{p}}; \mathbf{w}; x)}[\Pi_i^{\text{SWM}}(\hat{\mathbf{p}}; \mathbf{w}; x)] = \Pi_i^{\text{WS}}(\hat{\mathbf{p}}; \mathbf{w}; x)$. Because the deterministic WSWM satisfies properties ((a)-(f)) (Lambert et al. 2008), SWM also satisfies these properties. A realization of the score $\varphi \circ s_{\tilde{x}_i}(p_i)$ can be larger than 1, implying that agent i can lose (or win) more than what he can lose (or win) in the deterministic WSWM. However, we also notice that for some extreme values of error rates, the constraint $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) \geq -w_i$ can be violated⁵, i.e., an agent may lose more than his wager, which makes SWM invalid. In the next section, we show that by selecting error rates in a subtle way, we can obtain all the properties (a)-(h) without violating the wager constraint $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) \geq -w_i$.

SWM with Error rate selection (SWME) and random partition SWME (RP-SWME)

We notice that according to Lemma 1, no matter which error rates e_0, e_1 are chosen, the unbiasedness property of SWM holds, i.e., $\mathbb{E}_{\Pi_i \sim \mathcal{D}(\hat{\mathbf{p}}; \mathbf{w}; x)}[\Pi_i^{\text{SWM}}(\hat{\mathbf{p}}; \mathbf{w}; x)] = \Pi_i^{\text{WSWM}}(\hat{\mathbf{p}}; \mathbf{w}; x)$. In other words, we can choose the error rates in an arbitrary way (even depending on $\hat{\mathbf{p}}, \mathbf{w}$) without changing the expected net-payoff⁶ of each agent under any realized event outcome. This gives us the flexibility to tune the maximum amount of money each agent can win or lose in the game, while preserving the properties ((a)-(f)) inherited from WSWM.

Given reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} but not the event outcome x , the error rate pair that guarantees no wager violation under any outcome $x \in \mathcal{X}$ and any realization of the randomness induced by SWM may not be unique. We propose Algorithm 3 to select a pair of error rates e_0, e_1 after the reports and wagers are collected, such that at least one agent loses all his wager in the worst case w.r.t. the outcome and the randomness of SWM. We name the mechanism as SWME when we use Algorithm 3 to select the error rates for SWM.

Lemma 2. *SWME has no wager violation and when there exists at least one report $\hat{p}_i \neq 0.5$, at least one of the agents*

⁵For example, in a wagering game, two agents both wager 1 and report 1 and 0, respectively. Let $s_x(\hat{p}_i) = 1 - (x - \hat{p}_i)^2, e_j^i = 0.4, i = 1, 2, j = 0, 1$. In the worst case of agent 1, the surrogate outcomes are realized as $\tilde{x}_1 = 0, \tilde{x}_2 = 1$. Then, $\pi_1 = -5 < -1$.

⁶The expectation is taken over the randomness of the mechanism conditioned on the event outcome.

Algorithm 3 Error Rate Selection Algorithm

- 1: Collect the predictions $\hat{\mathbf{p}}$ and wagers \mathbf{w} .
 - 2: $\forall i: s_i^w \leftarrow \min_{x \in \mathcal{X}} s_x(\hat{p}_i), s_i^b \leftarrow \max_{x \in \mathcal{X}} s_x(\hat{p}_i)$.
 - 3: For each agent $i \in \mathcal{N}$, compute $r_i: r_i \leftarrow \frac{1}{2} + \frac{(1 - \frac{w_i}{W_{\mathcal{N}}})(s_i^w - s_i^b) + \sum_{j \in \mathcal{N} \setminus \{i\}} \frac{w_j}{W_{\mathcal{N}}}(s_j^w - s_j^b)}{2(2 + s_i^w + s_i^b - \sum_{j \in \mathcal{N}} \frac{w_j}{W_{\mathcal{N}}}(s_j^w + s_j^b))}$
 - 4: If $\min_{j \in \mathcal{N}} \{r_j\} = 0.5$, set $e_1^i = e_0^i = 0, \forall i$, else set $e_1^i = e_0^i = \min_{j \in \mathcal{N}} \{r_j\}, \forall i$.
-

Mechanism 4 Random Partition SWME (RP-SWME)

- 1: Partition agents into groups of two. If N is odd, leave one group with three agents.
 - 2: Run SWME for each group.
-

loses all his wager in the worst case w.r.t. the event outcome and the randomness of SWME.

Note Lemma 2 does not imply PO for SWME - if there exist two agents who have different predictions and have wager left even in their own worst cases, they can form a profitable bet against each other. We propose a variant of SWME to fix this caveat as follows.

Random partition SWME (RP-SWME)

Lemma 2 implies that when agents are partitioned into groups of two, there will not exist side bets. Meanwhile, a smaller number of agents imposes less restrictions in selecting the error rates, and thus each agent's wager can be fully leveraged in the randomization step. We would like to note that this is a very unique property of SWME: as both shown in Freeman, Pennock, and Vaughan (2017) and our experimental results, when the number of agents is small, existing wagering mechanisms (including DCA) all have low risk, i.e., have only a small portion of wager to lose in the worst case. This not only implies that SWME is particularly suitable for small group wagering but also points out a way to further improve the risk property of SWME, i.e. via randomly partitioning agents into smaller groups. We formally present the random partition SWME in Mechanism 4. We show in next section that RP-SWME achieves all properties (a)-(h).

Properties of SWME and RP-SWME

Theorem 2. *Both SWME and RP-SWME satisfy properties (a)-(g). In addition, RP-SWME satisfies (h).*

Proof. We provide full proofs in Section 6.3 of our full version (Chen, Liu, and Wang 2018), but we give the arguments for establishing surrogate wager's ex-post budget balance (despite of the randomness), incentive compatibility, and Pareto optimality.

Budget balance This can be shown via writing down the sum of net-payoffs defined in Eqn. (3). Our note below Eqn. (1) also states that the budget balance property doesn't depend on the specific forms of the scoring functions used.

Strict incentive compatibility First consider SWME. For an arbitrary profile of reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , Algorithm 3 outputs a profile \mathcal{E} of error rates of all agents. Denote by $\hat{\varphi}_{\mathcal{E}}^i(\cdot)$ the corresponding surrogate function specified using the error rate profile \mathcal{E} for agent i . For each i and $j \in \mathcal{N}$:

$$\begin{aligned} \mathbb{E}_{X \sim p_i, \tilde{X}_j} [\hat{\varphi}_{\mathcal{E}}^j \circ s_{\tilde{X}_j}(\hat{p}_j)] &= p_i \mathbb{E}_{\tilde{X}_j | X=1} [\hat{\varphi}_{\mathcal{E}}^j \circ s_{\tilde{X}_j}(\hat{p}_j)] \\ &+ (1 - p_i) \mathbb{E}_{\tilde{X}_j | X=0} [\hat{\varphi}_{\mathcal{E}}^j \circ s_{\tilde{X}_j}(\hat{p}_j)] \\ &= p_i \cdot s_{X=1}(\hat{p}_j) + (1 - p_i) \cdot s_{X=0}(\hat{p}_j) = \mathbb{E}_{X \sim p_i} [s_X(\hat{p}_j)], \end{aligned}$$

using Lemma 1. Then, using the linearity of expectation, we have (here $\tilde{\mathbf{X}}$ encodes the randomness in Π_i^{SWME})

$$\begin{aligned} &\mathbb{E}_{X \sim p_i, \tilde{\mathbf{X}}} [\Pi_i^{\text{SWME}}(\hat{\mathbf{p}}, \mathbf{w}, X)] \\ &= \frac{w_i W_{\mathcal{N} \setminus \{i\}}}{W_{\mathcal{N}}} \left(\mathbb{E}_{X \sim p_i, \tilde{X}_i} [\hat{\varphi}_{\mathcal{E}}^i \circ s_{\tilde{X}_i}(\hat{p}_i)] \right. \\ &\quad \left. - \sum_{j \in \mathcal{N} \setminus \{i\}} \frac{w_j}{W_{\mathcal{N} \setminus \{i\}}} \mathbb{E}_{X \sim p_i, \tilde{X}_j} [\hat{\varphi}_{\mathcal{E}}^j \circ s_{\tilde{X}_j}(\hat{p}_j)] \right) \\ &= \mathbb{E}_{X \sim p_i} \left[\frac{w_i W_{\mathcal{N} \setminus \{i\}}}{W_{\mathcal{N}}} \left(s_X(\hat{p}_i) - \sum_{j \in \mathcal{N} \setminus \{i\}} \frac{w_j}{W_{\mathcal{N} \setminus \{i\}}} s_X(\hat{p}_j) \right) \right] \\ &= \mathbb{E}_{X \sim p_i} [\Pi_i^{\text{WS}}(\hat{\mathbf{p}}, \mathbf{w}, X)]. \end{aligned}$$

Note the above holds for any possible reports ($\forall \mathcal{E}$). Thus the incentive properties of WSWM will preserve, i.e., truthfully reporting agent i 's private belief returns a higher payoff. The proof for RP-SWME is similar, with the only difference in that each agent's net-payoff is further averaged over the random partitions (but SIC under each possible partition).

Pareto optimality In RP-SWME, any pair of agents with different beliefs have a positive probability to be partitioned into a sub-group. Applying Lemma 2, at least one of them loses all his wager in the worst case. Thus, by Definition 7, RP-SWME is PO. \square

Wager with noisy ground truth

The above method also points out a way to implement a wagering mechanism with a noisy ground truth, as SWM is able to remove the noise in outcomes in expectation. This makes wagering possible even when only a noisy copy of outcome is available. Due to space limitation, we present the key idea below, while not re-defining all properties w.r.t. the noisy ground truth \hat{X} rather than X . The necessary changes are rather straight-forward.

Suppose we know a noisy estimate \hat{X} on X , and denote the error rates of \hat{X} as \hat{e}_1, \hat{e}_0 (which we know, and agents trust us in knowing these two numbers), we will be able to reproduce our surrogate wagering mechanism by plugging $\hat{X}, \hat{e}_1, \hat{e}_0$ into Eqn. (2), if we ignore the PO property for now. Similarly, we will have the wager violation issue pointed out earlier - we, however, do not have the control of the error rates directly. An easy fix is via the following affine transformation of the net-payoffs: given the error rates, suppose, over all possible predictions, wagers and the randomness of the mechanism, the largest wager-lose factor of an

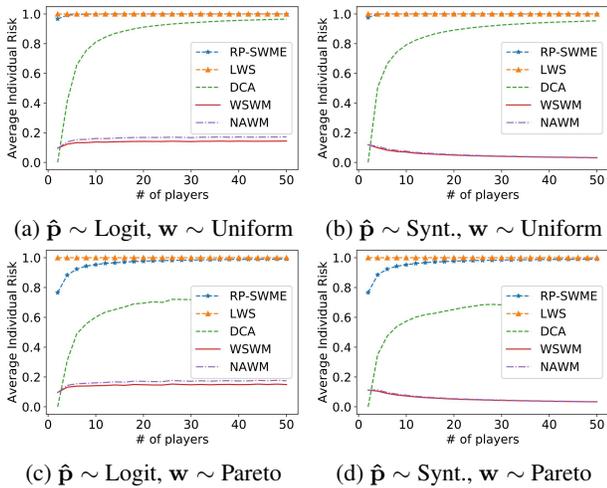


Figure 1: Average individual risks varying wagering mechanisms, prediction and wager distributions and # of agents N

agent that can be achieved is *scale*, i.e., an agent may lose at most $-\text{scale} \cdot w_i$ with $\text{scale} > 0$. We can then re-scale every agent’s net-payoff by $1/\text{scale}$. As this affine transformation is predetermined according to the worst case of all possible inputs, it does not affect the incentive and other properties of the original surrogate wagering mechanism, as $\mathbb{E}[\varphi \circ \Pi_i^{\text{WS}}(\cdot)] = \frac{1}{\text{scale}} \cdot \mathbb{E}[\Pi_i^{\text{WS}}(\cdot)]$.⁷ To achieve PO, we can further flip on \tilde{X} . To see this, denote the flipping error rates of $\tilde{X}|X$ as e_1, e_0 . Then the error rates of $\tilde{X}|X$ is given by (the deduction is given in Section 6.4 of the full version (Chen, Liu, and Wang 2018))

$$\begin{aligned} \mathbb{P}(\tilde{X}_i = 1|X = 0) &= e_0^i \cdot (1 - \hat{e}_0) + (1 - e_1^i) \cdot \hat{e}_0 \\ \mathbb{P}(\tilde{X}_i = 0|X = 1) &= e_1^i \cdot (1 - \hat{e}_1) + (1 - e_0^i) \cdot \hat{e}_1 \end{aligned}$$

It’s easy to see that when $\hat{e}_1 + \hat{e}_0 \neq 1$, we can tune the error rates of \tilde{X} via tuning e_1^i, e_0^i . This step corresponds to the error selection step in SWME.

6 Evaluation

In this section, we conduct simulations to evaluate the average-case performance of wagering mechanisms. We show that LWS and RP-SWME are more efficient than existing deterministic (weakly) incentive compatible mechanisms WSWM, NAWM and DCA. Meanwhile, we show that RP-SWME has smaller payoff variance and higher probability of winning money than LWS.

In the simulations, we generate the predictions of agents using two models: i). the logit model (Satopää et al. 2014), ii). the synthetic model (Ranjan and Gneiting 2010; Allard, Comunian, and Renard 2012; Satopää et al. 2014). The details of these models can be found in Section 8.1 of the full version (Chen, Liu, and Wang 2018). We generate the wagers of agents by two distributions: i). a uniform distribu-

⁷We didn’t apply the scaling in SWME, as the scaling will effectively decrease the expected payoff of each agent.

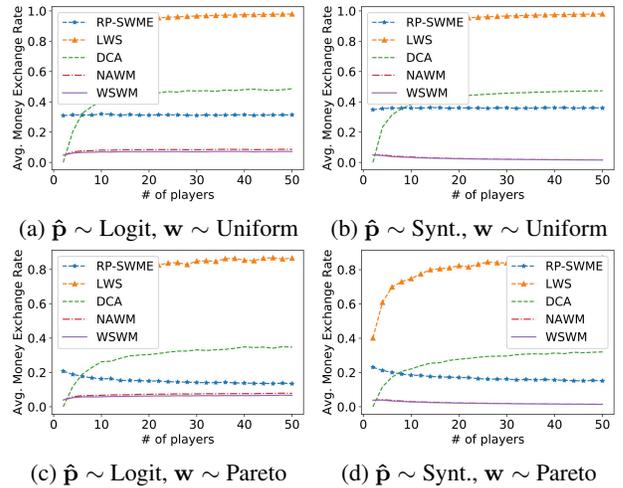


Figure 2: Average money exchange rate w.r.t. mechanisms, distributions of predictions and wagers, and # of agents N

tion over $[0,1]$, ii). a Pareto distribution with shape parameter 1.16 and scale parameter 1, characterizing the “20% of the population has 80% of the wealth” (Freeman, Pennock, and Vaughan 2017). We use Brier score as the scoring rule used in the wagering mechanisms that we evaluate.

Individual risk and money exchange rate. Individual risk and money exchange rate are two indicators of efficiency of wagering mechanisms.

Individual risk is the percent of wager that an individual agent can lose in the worst case w.r.t. the event outcome and the randomness of the mechanisms. The average individual risk is an indicator of Pareto optimality, because the average individual risk equaling to 1 (i.e., no one can commit to secure a positive wager before the wagering game) is a sufficient condition of Pareto optimality. *Money exchange rate* is the total amount of money exchanged in the game after the outcome of a wagering mechanism is realized, divided by the total amount of wagers. Average money exchange rate measures the efficiency of an average wagering game.

The simulations on binary events show that for both individual risk (Fig. 1) and money exchange rate (Fig. 2), RP-SWME and LWS outperform the incentive-compatible deterministic wagering mechanisms WSWM and NAWM. Moreover, LWS beats DCA, which was designed to increase the risk but is only weakly incentive compatible. The evaluations on multiple events show similar results.

Variance of payoff and probability of winning money. We evaluate these two metrics by considering an average agent with a prediction at different level of accuracy. The accuracy is measured by $1 - |x - \hat{p}_i|$. The predictions are generated according to a uniform distribution. Our simulations (Fig. 3) show that RP-SWME has smaller payoff variance and higher probability of winning money than LWS at each accuracy level.

More detailed evaluations can be found in Section 8 of the full version (Chen, Liu, and Wang 2018).

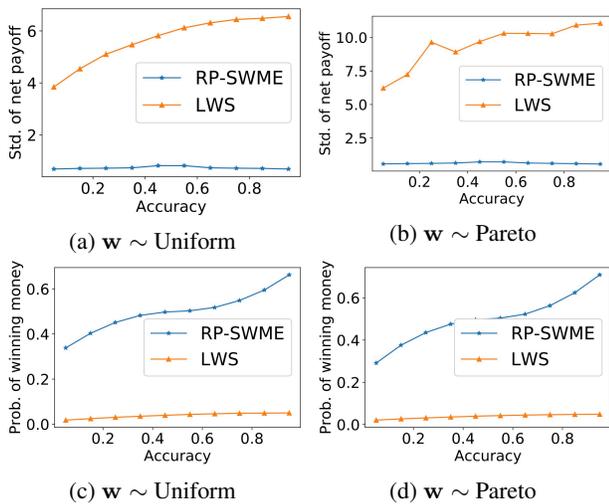


Figure 3: Variance of payoff and probability of winning money of RP-SWME and LWS

7 Conclusion

We extend the design of wagering mechanism to its randomized space. We propose two of them: Lottery Wagering Mechanisms (LWM) and Surrogate Wagering Mechanisms (SWM). We demonstrate the power of randomness by theoretically proving that they both satisfy a set of desirable properties, including Pareto optimality. SWM is also robust to noisy outcomes. We then carried out extensive experiments to support our theoretical findings.

References

Allard, D.; Comunian, A.; and Renard, P. 2012. Probability aggregation methods in geoscience. *Mathematical Geosciences* 44(5):545–581.

Brier, G. W. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78(1):1–3.

Bylander, T. 1994. Learning linear threshold functions in the presence of classification noise. In *Proceedings of the seventh annual conference on Computational learning theory*, 340–347. ACM.

Chen, Y.; Devanur, N. R.; Pennock, D. M.; and Vaughan, J. W. 2014. Removing arbitrage from wagering mechanisms. In *Proceedings of the 15th ACM EC*, 377–394. ACM.

Chen, Y.; Liu, Y.; and Wang, J. 2018. Randomized wagering mechanisms. *arXiv preprint arXiv:1809.04136*.

Cummings, R.; Pennock, D. M.; and Vaughan, J. W. 2016. The possibilities and limitations of private prediction markets. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 143–160. ACM.

Freeman, R., and Pennock, D. M. 2018. An axiomatic view of the parimutuel consensus wagering mechanism. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 1936–1938. International Foundation for Autonomous Agents and Multiagent Systems.

Freeman, R.; Pennock, D. M.; and Vaughan, J. W. 2017. The double clinching auction for wagering. In *Proceedings of the 18th ACM EC*, 43–60. ACM.

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

Johnstone, D. J. 2007. The parimutuel Kelly probability scoring rule. *Decision Analysis* 4(2):66–75.

Jose, V. R.; Nau, R. F.; and Winkler, R. L. 2006. Scoring rules, generalized entropy and utility maximization. Working Paper, Fuqua School of Business, Duke University.

Kilgour, D. M., and Gerchak, Y. 2004. Elicitation of probabilities using competitive scoring rules. *Decision Analysis* 1(2):108–113.

Lambert, N. S.; Langford, J.; Wortman, J.; Chen, Y.; Reeves, D.; Shoham, Y.; et al. 2008. Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM EC*, 170–179. ACM.

Lambert, N. S.; Langford, J.; Vaughan, J. W.; Chen, Y.; Reeves, D. M.; Shoham, Y.; and Pennock, D. M. 2015. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory* 156:389–416.

Liu, Y., and Chen, Y. 2018. Surrogate Scoring Rules and a Dominant Truth Serum for Information Elicitation. *arxiv preprint*.

Matheson, J. E., and Winkler, R. L. 1976. Scoring rules for continuous probability distributions. *Management Science* 22(10):1087–1096.

Natarajan, N.; Dhillon, I. S.; Ravikumar, P. K.; and Tewari, A. 2013. Learning with noisy labels. In *Advances in neural information processing systems*, 1196–1204.

Ranjan, R., and Gneiting, T. 2010. Combining probability forecasts. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(1):71–91.

Satopää, V. A.; Baron, J.; Foster, D. P.; Mellers, B. A.; Tetlock, P. E.; and Ungar, L. H. 2014. Combining multiple probability predictions using a simple logit model. *International Journal of Forecasting* 30(2):344–356.

Scott, C. 2015. A rate of convergence for mixture proportion estimation, with application to learning from noisy labels. In *AISTATS*.

Winkler, R. L. 1969. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association* 64(327):1073–1078.

Witkowski, J.; Freeman, R.; Vaughan, J. W.; Pennock, D. M.; and Krause, A. 2018. Incentive-compatible forecasting competitions. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*.