

An Axiomatic View of the Parimutuel Consensus Wagering Mechanism

Rupert Freeman
 Duke University
 rupert@cs.duke.edu

David M. Pennock
 Microsoft Research
 dpennock@microsoft.com

Abstract

We consider an axiomatic view of the *Parimutuel Consensus Mechanism (PCM)* defined by Eisenberg and Gale [1959]. The parimutuel consensus mechanism can be interpreted as a parimutuel market for wagering with a proxy that bets optimally on behalf of the agents, depending on the bets of the other agents. We show that the parimutuel consensus mechanism uniquely satisfies the desirable properties of Pareto optimality, individual rationality, budget balance, anonymity, sybilproofness and envy-freeness. While the parimutuel consensus mechanism does violate the key property of incentive compatibility, it is incentive compatible in the limit as the number of agents becomes large. Via simulations on real contest data, we show that violations of incentive compatibility are both rare and only minimally beneficial for the participants. This suggests that the parimutuel consensus mechanism is a reasonable mechanism for eliciting information in practice.

1 Introduction

In 1867, Spanish entrepreneur Joseph Oller invented *parimutuel betting*, a form of wagering still popular today, handling billions of dollars annually on horse races and jai alai games. Each bettor places money on one of several future outcomes—say, horse #1 to win a race. She is allowed to cancel her bet or move her money to a different outcome at any time, even at the last second before the race begins. After the outcome resolves—say, horse #1 wins—agents who picked the wrong outcome lose their wagers to the agents who picked correctly. Winning agents split the pot in proportion to the size of their wagers.

Eisenberg and Gale [1959] analyzed the equilibrium of parimutuel betting, defining the *parimutuel consensus mechanism (PCM)*. The PCM is equivalent to parimutuel betting where each agent has a proxy. Each agent’s proxy knows her true probabilities for all outcomes. As bets come in, and the prospective payoff per dollar, or *odds*, for each outcome converge, the proxy automatically switches its agent’s money to the outcome yielding the highest expected payoff for that

agent. In equilibrium, all the proxies are optimizing and none want to switch outcomes.

At any point in time, the odds can be interpreted as probabilities, providing a prediction of the outcome of the event. Thus, facilitating wagering can serve as a source of free information for a principal seeking to forecast a future event.

Eisenberg and Gale discuss one undesirable feature of the equilibrium: it produces odds that sometimes ignore some agents. Manski [2006] further explores in detail how the equilibrium of risk-neutral, budget-constrained agents may fail to aggregate beliefs in a sensible way. Additionally, the PCM is not *incentive compatible*, or truthful. An agent may strategically improve her payoff by taking into account what other agents know or what they may do. In the end, her best action may be to report false probabilities to her proxy that differ from her true subjective probabilities. For a principal whose primary goal is information elicitation, this is problematic because some of the reported probabilities may not faithfully reflect the bettors’ private information.

Given the potential for bad equilibria and the lack of incentive compatibility, why is the PCM still prevalent? One answer is that, in practice, it often works fine. Parimutuel betting does consistently induce a wisdom-of-crowds effect, producing odds that encode well calibrated and accurate probabilistic forecasts of the outcomes [Ali, 1977; Snyder, 1978; Thaler and Ziemba, 1988], like many prediction markets do [Arrow *et al.*, 2008]. Plott *et al.* [2003] tested parimutuel betting in a laboratory experiment, showing that the mechanism is an effective vehicle for information aggregation regardless of why it might go wrong in theory. If agents have concave or risk-averse utility for money, the equilibrium of similar mechanisms is stable and induces sensible belief aggregation [Wolfers and Zitzewitz, 2006; Beygelzimer *et al.*, 2012]. In particular, an agent with logarithmic utility does best by betting an amount on each outcome proportional to her probability [Cover and Thomas, 2006].

We examine another plausible reason why the PCM continues to enjoy usage: the mechanism satisfies a number of desirable axioms for wagering systems. We prove that the PCM is the unique wagering mechanism that is Pareto optimal, individually rational, budget balanced, sybilproof, anonymous, and envy-free, subject to a mild condition on the reports.

Freeman *et al.* [2017] show that, to gain the key property of incentive compatibility, one of the core properties must be

relaxed. Yet we can show that the PCM is near incentive compatible in some cases. Yes, there are scenarios where agents can gain from lying, but we prove that the PCM is incentive compatible *in the large*, as the number of agents grows. In extensive simulations using real forecasts from an online contest, we show that opportunities for agents to profit from untruthful play are rare, mostly vanishing as the number of agents grows. Our results shed light on the practical success of the PCM. Despite its flaws, identified as early as 1959, it does satisfy six natural and desirable properties of wagering mechanisms and it comes close both theoretically and empirically to obtaining a crucial seventh: incentive compatibility.

2 Preliminaries

Consider a random variable (*event*) X which takes a value (*outcome*) in $\{0, 1\}$. There is a set of agents (or bettors) \mathcal{N} , each with a private, subjective, immutable belief p_i which is the probability that $X = 1$, and a budget w_i , which is the maximum amount that they are prepared to lose.

A *wagering mechanism* is used to elicit beliefs from the agents. Each agent submits a report $\hat{p}_i \in [0, 1]$ and wager $w_i \in \mathbb{Q}_{\geq 0}$ to the mechanism, where \hat{p}_i captures her belief and w_i captures her budget.¹ We denote the vector of agents' reports by $\hat{\mathbf{p}}$, and the vector of agents' wagers by \mathbf{w} . The wagering mechanism defines a payoff $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x)$ to each agent that depends on the reports and wagers of the agents and x , the observed value of X . To be a valid wagering mechanism, it must be the case that no agent loses more than her wager (i.e., $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) \geq -w_i$) and that an agent with zero wager does not participate in the mechanism (i.e., $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) = 0$ if $w_i = 0$). We can therefore let $\mathcal{N} = \mathbb{N}$ without loss of generality, because non-participation is equivalent to a zero wager. We denote the reports and wagers of all agents other than i by $\hat{\mathbf{p}}_{-i}$ and \mathbf{w}_{-i} respectively.

2.1 Security Interpretation

Freeman *et al.* [2017] observed that the output of a wagering mechanism can be interpreted as an allocation of Arrow-Debreu securities with payoffs that depend on the realization of X . A *yes* security is a contract that pays off \$1 in the outcome $X = 1$ and \$0 if $X = 0$. Similarly, a *no* security pays off \$0 if $X = 1$ and \$1 if $X = 0$. A *risk neutral* agent with belief p about the likelihood that $X = 1$ would be willing to buy a *yes* security at any price up to p or a *no* security at any price up to $1 - p$.

Suppose a wagering mechanism would yield a net payoff to agent i of $\rho_1 = \Pi_i(\hat{\mathbf{p}}; \mathbf{w}; 1)$ when $X = 1$ and $\rho_0 = \Pi_i(\hat{\mathbf{p}}; \mathbf{w}; 0)$ when $X = 0$. This is equivalent to the payoff that i would receive if she were sold $y_i = \max\{\rho_1 - \rho_0, 0\}$ *yes* securities and $n_i = \max\{\rho_0 - \rho_1, 0\}$ *no* securities for a total cost of $\sigma_i = \max\{-\rho_0, -\rho_1\}$. For example, if $\rho_0 < \rho_1$, then agent i 's participation in the wagering mechanism is equivalent to agent i paying the principal $\sigma_i = -\rho_0$ before X is realized and then receiving $y_i = \rho_1 - \rho_0$ from the principal in the outcome $X = 1$.

¹The assumption of rational wagers is required in Section 4. Rational wagers can approximate real-valued budgets arbitrarily well.

Therefore, the output of a wagering mechanism can be completely specified by a triple $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma})$, where for each agent i , y_i (resp. n_i) is the number of *yes* (*no*) securities allocated to i , and σ_i is the cost paid by i for these securities. The requirement that no agent can lose more than her wager is captured by the constraint $\sigma_i \leq w_i$. Without loss of generality, we assume that for all i , either $y_i = 0$ or $n_i = 0$ (or both), since any (fraction of a) pair of *yes* and *no* securities can be precisely converted into (a fraction of) \$1.

2.2 Properties of Wagering Mechanisms

Lambert *et al.* [2008] introduced several desirable properties for wagering mechanisms. We focus on five of these properties in our analysis: anonymity, individual rationality, incentive compatibility, budget balance, and sybilproofness.

First, *anonymity* says that the payouts do not depend on the identities of the agents. This is a basic property that all wagering mechanisms proposed in the literature, including the PCM, satisfy.

Individual rationality says that agents do not lose money (in expectation) by participating truthfully. We note that this is slightly stronger than the usual definition, which only requires that there exists *some* (not necessarily truthful) report for which agent i does not make an expected loss. Our definition is implied by the combination of the usual definition of individual rationality, and incentive compatibility. We require the stronger version in the proof of Theorem 2.

Definition 1. A *wagering mechanism* is individually rational if, for any agent i and any belief p_i , for all $\hat{\mathbf{p}}_{-i}, \mathbf{w}$ her expected payoff is greater than or equal to her cost; that is,

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) \geq \sigma_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w})$$

Incentive compatibility requires that each agent maximizes her expected payoff by reporting truthfully, regardless of the reports and wagers of other agents.

Definition 2. A *wagering mechanism* is (weakly) incentive compatible if, for every agent i with belief p_i and all reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} ,

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) - \sigma_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) \geq p_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - p_i) n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w}).$$

A wagering mechanism is budget balanced if the principal never makes a profit or a loss.

Definition 3. A *wagering mechanism* is budget balanced if, for all $\hat{\mathbf{p}}$ and \mathbf{w} , $\sum_{i \in \mathcal{N}} y_i(\hat{\mathbf{p}}; \mathbf{w}) = \sum_{i \in \mathcal{N}} n_i(\hat{\mathbf{p}}; \mathbf{w}) = \sum_{i \in \mathcal{N}} \sigma_i(\hat{\mathbf{p}}; \mathbf{w})$

If the mechanism sometimes makes a profit, but never a loss, then we say that it is *weakly budget balanced*.

A mechanism is sybilproof if it is not beneficial for agents to participate under multiple fake identities, or for agents reporting the same probability to merge.

Definition 4. A *wagering mechanism* is sybilproof if for any subset of players S , for any $\hat{\mathbf{p}}$ with $\hat{p}_i = \hat{p}_j$ for $i, j \in S$, for any vectors of wagers \mathbf{w}, \mathbf{w}' with $w_i = w'_i$ for $i \notin S$ and

$\sum_{i \in S} w_i = \sum_{i \in S} w'_i$, it is the case that:

$$\begin{aligned} & \sum_{i \in S} (y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) \\ &= \sum_{i \in S} (y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')) \end{aligned}$$

and for all $i \notin S$,

$$\begin{aligned} & (y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) \\ &= (y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')). \end{aligned}$$

Freeman *et al.* [2017] introduce the notion of *Pareto optimality* for wagering mechanisms, which is satisfied if there is no side bet that can be made by the agents on top of that already facilitated by the wagering mechanism, such that some agent benefits without harming others. Freeman *et al.* show that the following condition is equivalent.

Definition 5. A wagering mechanism is *Pareto optimal* if, for all reports $\hat{\mathbf{p}}$ and \mathbf{w} , there exists a $p \in [0, 1]$ such that

$$\begin{aligned} \forall i : \hat{p}_i < p, & \quad \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i \quad \text{and} \quad y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0, \\ \forall i : \hat{p}_i > p, & \quad \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i \quad \text{and} \quad n_i(\hat{\mathbf{p}}; \mathbf{w}) = 0. \end{aligned}$$

Freeman *et al.* [2017] show that four key properties are incompatible.

Theorem 1 (Freeman *et al.* [2017]). *No wagering mechanism simultaneously satisfies individual rationality, weak incentive compatibility, weak budget balance, and Pareto optimality.*

Lastly, we consider the property of envy-freeness [Foley, 1967]. Envy-freeness is a basic fairness property which says that no player should envy the allocation of securities to any other agent.

Definition 6. Say that agent i envies another agent j if $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) \leq w_i$ and

$$\begin{aligned} & \hat{p}_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i) n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) \\ & < \hat{p}_j y_j(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_j) n_j(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_j(\hat{\mathbf{p}}; \mathbf{w}) \end{aligned}$$

A wagering mechanism is *envy-free* if there is no pair of agents (i, j) such that i envies j .

Proofs are omitted or heavily abbreviated, but appear in the full version of the paper on the authors' websites.

3 The Parimutuel Consensus Mechanism

The Parimutuel Consensus Mechanism (PCM) can be thought of as a direct implementation of the equilibrium of parimutuel betting. The PCM includes the rules of parimutuel betting plus, conceptually, a proxy agent that automatically switches its agent's bet to the outcome with highest expected profit per security. The output of the mechanism is the equilibrium where all proxies are stable. For the binary case of `yes` and `no` outcomes that we consider, the PCM is defined by a price π such that an agent with report less than π is allocated `no` securities at a price of $1 - \pi$ per security, and an agent with report more than π is allocated `yes` securities at a price of π per security. The equilibrium condition is

$$\sum_{i: \hat{p}_i < \pi} \frac{w_i}{1 - \pi} + \sum_{i: \hat{p}_i = \pi} \frac{c_1 w_i}{1 - \pi} = \sum_{i: \hat{p}_i > \pi} \frac{w_i}{\pi} + \sum_{i: \hat{p}_i = \pi} \frac{c_2 w_i}{\pi}, \quad (1)$$

where c_1 and c_2 lie in the interval $[0, 1]$ and $\min\{c_1, c_2\} = 0$. These represent the fractions of budgets spent by agents with $\hat{p}_i = \pi$ who bet (some of) their wager to correctly balance the market prices and allow the market to reach equilibrium, even though they get zero expected profit. At most one of c_1 and c_2 is greater than 0, since it would be redundant to have agents with $\hat{p}_i = \pi$ betting on both `yes` and `no`. Note that the left hand side of Equation 1 is the total number of `no` securities allocated, and the right hand side is the total number of `yes` securities allocated. Eisenberg and Gale [1959] show as their main contribution that such a price is both unique and guaranteed to exist. The output of the PCM is defined by

$$y_i(\hat{\mathbf{p}}; \mathbf{w}) = \begin{cases} 0 & \hat{p}_i < \pi \\ \frac{c_2 w_i}{\pi} & \hat{p}_i = \pi \\ \frac{w_i}{\pi} & \hat{p}_i > \pi \end{cases}, \quad n_i(\hat{\mathbf{p}}; \mathbf{w}) = \begin{cases} \frac{w_i}{1 - \pi} & \hat{p}_i < \pi \\ \frac{c_1 w_i}{1 - \pi} & \hat{p}_i = \pi \\ 0 & \hat{p}_i > \pi \end{cases}$$

and

$$\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = \begin{cases} w_i & \hat{p}_i < \pi \\ \max\{c_1, c_2\} w_i & \hat{p}_i = \pi \\ w_i & \hat{p}_i > \pi \end{cases}$$

Example 1. Suppose that there are four agents, with reports $\hat{\mathbf{p}} = (0.3, 0.5, 0.6, 0.8)$ and wagers $\mathbf{w} = (1, 1, 3, 6)$. Observe that setting $\pi = 0.6$ and $c_1 = 2/3, c_2 = 0$ satisfies Equation 1: each side of the equation has value 10. Further, setting $\pi < 0.6$ results in the right hand side of Equation 1 being greater than the left hand side, for any allowed values of c_1 and c_2 , and the opposite is true for any $\pi > 0.6$.

We can now compute the output of the PCM on this instance, according to the formulae above. Agents 1 and 2 are allocated 2.5 `no` securities each, for a price of 1, agent 3 is allocated 5 `no` securities for a price of 2 (note that this is a $c_1 = 2/3$ fraction of agent 3's wager), and agent 4 is allocated 10 `yes` securities for a price of 6.

Recall that, by Theorem 1, no wagering mechanism can simultaneously satisfy individual rationality, incentive compatibility, weak budget balance, and Pareto optimality. Theoretical papers on wagering mechanisms are generally reluctant to give up any of the first three properties, sacrificing Pareto optimality [Lambert *et al.*, 2008; 2015; Chen *et al.*, 2014; Freeman *et al.*, 2017]. However, in practice, Pareto optimality is an important consideration and virtually all real-world wagering mechanisms, including parimutuels, bookmakers, and double auctions, satisfy it. This is because trade drives participation; a mechanism that facilitates little trade is of little use or interest to agents.

Individual rationality seems hard to give up. We cannot force agents to play a game that they expect to lose and, even if we did, they could just wager $w_i = 0$. The center may be willing to pay for the information inherent in the agents' beliefs, subsidizing the mechanism and relaxing budget balance. Market scoring rules [Hanson, 2003; Chen and Pennock, 2007], for example, do just that, losing a strictly bounded amount of money in service of gaining information. However, a patron will only subsidize events that bear on valuable decisions. Nearly all fielded wagering mechanisms have taxes, not subsidies, yielding profits, not losses.

If we want Pareto optimality, individual rationality, and (weak) budget balance, we are forced to give up on incentive compatibility. That's exactly what the PCM does. In the remainder of this paper, we show that the PCM is the unique wagering mechanism that simultaneously satisfies the other six properties from Section 2, subject to a condition on the reports. We then show that, despite not satisfying incentive compatibility, the PCM is approximately incentive compatible in two senses. First, we prove that, as the number of agents grows, the mechanism is incentive compatible *in the large*. Second, we show empirically that, across thousands of simulated wagers based on real probability estimates, opportunities to profit from misreports are almost negligible.

4 Properties of the PCM

Despite its theoretical flaws, including the possibility of non-sensical information aggregation, the PCM seems well behaved in practice. In this section, we examine one possible reason for this by theoretically justifying the PCM. We first note that the PCM satisfies six desirable properties for wagering mechanisms. Although incentive compatibility is not one of the six, we know that adding it is impossible without giving something up: no mechanism satisfying even just the first three properties can also be incentive compatible.

Proposition 1. *The parimutuel consensus mechanism satisfies individual rationality, budget balance, Pareto optimality, anonymity, sybilproofness, and envy-freeness.*

That the PCM satisfies the first three properties is noted by Freeman *et al.* [2017].

4.1 Axiomatic Characterization

We next show that the PCM is the *unique* wagering mechanism satisfying the six properties from Proposition 1, subject to a condition on the reports. Suppose that some non-zero wager is placed on $N > 3$ distinct reports, denoted by $P_1 < P_2 < \dots < P_N$, and let $W_k = \sum_{i:\hat{p}_i=P_k} w_i$ be the total wager at probability P_k . We say that the *non-extreme* assumption holds if $P_2 W_1 < (1 - P_2) \sum_{i=3}^N W_i$ and $(1 - P_{N-1}) W_N < P_{N-1} \sum_{i=1}^{N-2} W_i$. For the data set used in Section 5 and wagers generated according to a Pareto($\alpha = 1.16$) distribution (see Section 5 for details), the non-extreme assumption held on over 99.97% of instances.

Theorem 2. *Let M be a wagering mechanism satisfying anonymity, individual rationality, budget balance, Pareto optimality, sybilproofness, and envy-freeness. If the non-extreme assumption holds, then payoffs defined by M match those defined by the parimutuel consensus mechanism.*

Proof Sketch. We first show that any wagering mechanism that satisfies envy-freeness, sybilproofness, and anonymity is defined by fixed prices p_y and p_n for yes and no securities. That is, for all agents i with $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$, we have $p_y = \frac{\sigma_i(\hat{\mathbf{p}}; \mathbf{w})}{y_i(\hat{\mathbf{p}}; \mathbf{w})}$, and for all agents i with $n_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$, we have $p_n = \frac{\sigma_i(\hat{\mathbf{p}}; \mathbf{w})}{n_i(\hat{\mathbf{p}}; \mathbf{w})}$.

To prove this, suppose otherwise for contradiction. That is, suppose that there exist agents i, j with $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$

and $y_j(\hat{\mathbf{p}}; \mathbf{w}) > 0$ such that $\frac{\sigma_i(\hat{\mathbf{p}}; \mathbf{w})}{y_i(\hat{\mathbf{p}}; \mathbf{w})} > \frac{\sigma_j(\hat{\mathbf{p}}; \mathbf{w})}{y_j(\hat{\mathbf{p}}; \mathbf{w})}$. Consider a modified instance $(\hat{\mathbf{p}}, \mathbf{w}')$ in which both of i and j participate as sybils, denoted by sets S_i and S_j , instead of their individual identities, such that for all $k, \ell \in S_i \cup S_j$, we have that $\sigma_k = \sigma_\ell$. By sybilproofness and anonymity it must be the case that $\sigma_k(\hat{\mathbf{p}}; \mathbf{w}') = \sigma_i(\hat{\mathbf{p}}; \mathbf{w})/|S_i|$ and $y_k(\hat{\mathbf{p}}; \mathbf{w}') = y_i(\hat{\mathbf{p}}; \mathbf{w})/|S_i|$ for all $k \in S_i$, with the equivalent equalities holding for all $\ell \in S_j$ also. Therefore, for all $k \in S_i$ and $\ell \in S_j$,

$$\frac{\sigma_k(\hat{\mathbf{p}}; \mathbf{w}')}{y_k(\hat{\mathbf{p}}; \mathbf{w}')} = \frac{\sigma_i(\hat{\mathbf{p}}; \mathbf{w})}{y_i(\hat{\mathbf{p}}; \mathbf{w})} > \frac{\sigma_j(\hat{\mathbf{p}}; \mathbf{w})}{y_j(\hat{\mathbf{p}}; \mathbf{w})} = \frac{\sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')}{y_\ell(\hat{\mathbf{p}}; \mathbf{w}')}$$

Because $\sigma_k(\hat{\mathbf{p}}; \mathbf{w}') = \sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')$, k envies ℓ , violating envy-freeness in the modified instance. An identical argument shows the existence of a fixed price p_n for no securities.

By budget balance, the wagering mechanism must sell exactly the same number of yes and no securities, and it must be the case that each yes/no pair sells for exactly \$1. Therefore, $p_y + p_n = 1$. By individual rationality, it must be the case that all agents with $\hat{p}_i < p_y$ have $y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$, and all agents with $\hat{p}_i > p_y$ have $n_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$.

Next, we can use Pareto optimality, combined with anonymity and sybilproofness, to show that $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$ for all i with $p_i \neq p_y$, provided that p_y does not lie between P_1 and P_2 or between P_{N-1} and P_N . The idea is that if some agent i (say with $p_i > p_y$) is not spending her full wager on yes securities, then Pareto optimality implies that all other agents are spending their full wagers. As long as one such agent is allocated yes securities, which is guaranteed by the condition on the value of p_y , then i envies that agent (again, by sybilproofness, we can create an instance where all agents have equal wagers). The condition on the value of p_y holds due to the non-extreme assumption, which says that agents at the highest (resp. lowest) price cannot buy enough yes (no) securities to balance the demand for no (yes) securities from all other agents.

Finally, it is easy to show that any prices other than $p_y = \pi$ and $p_n = 1 - \pi$ are ruled out by the previous two paragraphs. We can apply anonymity and sybilproofness to show that any difference in the total number of yes and no securities sold to those agents with $\hat{p}_i \neq p_y$ must be compensated for by selling securities to agents with $\hat{p}_i = p_y$, in proportion to their wagers, as the PCM does. \square

4.2 Incentive Properties of the PCM

As a Corollary of Theorem 1 and Proposition 1, we know that the PCM violates incentive compatibility. Intuitively, this is because agents are able to change the price π by changing their reports.

Example 2. *Let $\mathbf{p} = (0.4, \frac{2}{3}, 0.8)$ and $\mathbf{w} = (1, 1, 1)$. Then the outcome of the PCM is $(\mathbf{y} = (0, 1.5, 1.5), \mathbf{n} = (3, 0, 0), \boldsymbol{\sigma} = (1, 1, 1))$. Note that the price $\pi = \frac{2}{3}$, so agent 2's utility is 0. However, if agent 2 misreports $\hat{p}_2 = 0.6$, then the outcome becomes $(\mathbf{y} = (0, \frac{5}{6}, \frac{5}{3}), \mathbf{n} = (2.5, 0, 0), \boldsymbol{\sigma} = (1, 0.5, 1))$. Now the price $\pi = 0.6$, so agent 2's utility is $\frac{5}{6} \cdot \frac{2}{3} - 0.5 = \frac{1}{18} > 0$.*

Example 2 has a particular form common to all profitable misreports. In order to change the price in a profitable way,

a manipulating agent must ensure that her misreport exactly matches the new equilibrium price. The intuition is that the only way an agent can affect the price is to report a probability on the opposite side of the current price as her belief. However, such a misreport is only profitable if she does not ‘over-shoot’ and end up buying the wrong type of security.

Theorem 3. *Let $\hat{p}_i \neq p_i$ be a profitable misreport for agent i . Let π_T denote the *yes* security price when agent i reports truthfully, and π_M denote the *yes* security price in the instance when i misreports \hat{p}_i . Then it must be the case that $\hat{p}_i = \pi_M$, and either $\hat{p}_i < \pi_T \leq p_i$ or $p_i \leq \pi_T < \hat{p}_i$.*

Incentive Compatibility in the Large. We now show that the PCM satisfies an approximate notion of incentive compatibility known as *incentive compatibility in the large (IC-L)*, introduced by Azevedo and Budish [2017].² It relaxes incentive compatibility by requiring only that truthful reporting is optimal as the number of agents grows large, and that truthful reporting is only optimal in expectation over the reports, rather than based on the (ex-post) realization of reports, as our definition of incentive compatibility requires. We refer the reader to Azevedo and Budish [2017], or the full version of this paper, for a formal definition of IC-L.

Theorem 4. *The parimutuel consensus mechanism satisfies incentive compatibility in the large.*

In cases where only a finite set of reports are allowed, Theorem 4 follows directly from envy-freeness (Azevedo and Budish show that envy-freeness implies IC-L when the number of possible reports is finite). Since finite sets of reports can provide arbitrary precision, this is usually enough for practical purposes. Most real-world mechanisms only allow reports up to a precision of 1%, and this is also the precision we use in our simulations (see Footnote 4). However, for completeness, we also show that the PCM satisfies IC-L for the case where an infinite number of reports are allowed. The intuition for the result is that a single agent’s ability to influence the price π decreases as the number of agents increases.

5 Simulations

We tested the incentive compatibility of the PCM on a data set consisting of probability reports gathered from an online prediction contest called ProbabilitySports [Galebach, 2004].³ The data set consists of probabilistic predictions about the outcome of 1643 National Football League matches from the start of the 2000 NFL preseason until the end of the 2004 season. For each match, between 64 and 1574 players reported their subjective probability of a fixed team (say, the home team) winning the match. Each match was scored according to the Brier scoring rule, with points contributing to a season-long scoreboard.

ProbabilitySports users submitted probabilities but not wagers. We generated wagers from a variety of Pareto distributions.

Pareto distributions are a natural choice as they approximately model the distribution of wealth in a population. A Pareto distribution is defined by two parameters: a scale parameter k , which has the effect of multiplicatively scaling the distribution, and a shape parameter α , which affects the size of the distribution’s tail. To allow for a fair comparison between distributions and instance sizes, we scaled each set of randomly generated wagers so that the average wager is 1. This means that changing the scale parameter has no effect, as the wagers are rescaled anyway. Therefore, we fix the scale parameter to 1 and vary only the shape parameter.

The first Pareto distribution we use for wager generation has $\alpha = 1.16$, which is often described as “20% of the population has 80% of the wealth,” and classically viewed as a realistic distribution of wealth. Second, we use $\alpha = 3$, which produces a more equal distribution of wagers in comparison to $\alpha = 1.16$. Finally, we consider a uniform distribution of wagers (that is, $w_i = 1$ for all agents), corresponding to a situation either where all agents are equal, or where they do not have the opportunity to choose their wager (as in the ProbabilitySports competition). Note that the uniform distribution is the limit of the Pareto distribution as $\alpha \rightarrow \infty$.

Our first step was to examine the entire dataset. For each of the 1643 matches and each wager distribution, we randomly generated a set of wagers drawn from that distribution. For each set of wagers we chose 50 random agents and simulated 101 reports for them in the range $\{0, 0.01, \dots, 0.99, 1\}$.⁴ For each report, we computed the agent’s expected utility, taking their true belief to be their original report p_i . If there exists a misreport $\hat{p}_i \neq p_i$ that yields a higher utility than reporting their true belief, then the agent has a profitable misreport.

The results are summarized in Table 1. We report four statistics. The ‘% Agents With Profitable Misreports’ column states the percentage of agents that are able to benefit from misreporting. The ‘Average Profit’ column gives, out of those agents with a profitable misreport, the average benefit that the agent can gain from misreporting optimally, over and above her utility from reporting truthfully. The ‘Average Wager per Profitable Misreport’ column gives the average wager of agents with a profitable misreport available. Finally, the ‘Average Misreport Distance’ gives, for those agents with a profitable misreport, the average distance between the optimal misreport and the true belief.

For wagers generated from a Pareto distribution with $\alpha = 1.16$, we found 55 profitable misreports (out of 82,150 agents that we checked), which means that only around 0.07% of the agents that we checked had a profitable misreport. It is striking to consider the makeup of this small percentage of agents. The average wager of these agents is 118.8 (recall that each set of wagers is scaled so that the average wager is 1). What we are seeing is that only agents with very, very high wagers have sufficient power to change the price π . In contrast, the average profit that these agents obtain by misreporting is only


²There is a large body of work focusing on other limiting IC criteria, including ϵ -strategyproofness [Roberts and Postlewaite, 1976; Ehlers *et al.*, 2004], that we do not focus on here.

³We thank Brian Galebach for providing us with this data.

⁴In principle, our setup allows agents to report at a higher precision than this, so there will be some possible misreports that we do not detect. However, we believe that considering reports of only multiples of 0.01 is reasonable, due to limited cognitive capacity of the agents and the practical constraints of many wagering systems.

	% Agents With Profitable Misreports	Average Profit	Average Wager per Profitable Misreport	Average Misreport Distance
Pareto($\alpha = 1.16$)	0.07	1.55	118.8	0.044
Pareto($\alpha = 3$)	< 0.01	0.03	5.76	0.015
Uniform	0	n/a	n/a	n/a

Table 1: Profitable misreports under Pareto and uniform wager generation.


 Figure 1: Profitable misreport characteristics for varying wager distributions and varying numbers of agents, n .

1.55, suggesting that even these high-wager agents are unable to have too large an effect on the security price π . This average profit is on the order of 1-2% of the misreporting agents' wagers. For those agents that do misreport, the optimal misreport only differs from their belief by around 0.04.

As α increases, the number of agents with opportunity to misreport decreases. Indeed, for the uniform wagers, we did not find a single opportunity to profitably misreport. This is not surprising, since when wagers are uniform and there are a large number of agents, no agent will ever be able to significantly affect the price.

So for the full data set, with $64 \leq n \leq 1574$ agents per match, opportunities to profitably misreport are scarce, as we would expect because the PCM satisfies IC-L. But what about instances with fewer agents? To investigate smaller instances, we subsampled smaller values of n from the complete set of reports and ran the same simulation. For each match, each value of $n \in \{10, 20, 30, 40, 50\}$, and each wager distribution, we randomly sampled n reports and generated wagers. For every instance generated in this way, we tested *every* agent to see whether they had a profitable misreport.

Figure 1a shows how the percentage of agents that can profitably misreport changes with instance size. With only 10 agents per instance, around 10% of all agents have a profitable misreport available to them. This fraction decreases quickly as n increases – for instances with 50 agents, less than 2% of agents are able to profitably misreport. Interestingly, all wager distributions exhibit approximately the same susceptibility to manipulation, in contrast to the full instances. We speculate that this is because, while high-wager agents are more likely to have profitable manipulations available, their existence also prevents low-wager agents from being able to manipulate, thus rendering the existence of high-wager agents something of a wash for small n . For large n , the latter effect disappears, since low-wager agents are unable to profitably misreport, even in the absence of high-wager agents.

Figure 1b shows how the average value of each profitable misreport changes with n , where the value of a profitable misreport is the difference in expected utility between the agent's optimal misreport and their truthful report. Interestingly, we

see three very different trends depending on the wager distribution, all of which are consistent with the results on the full dataset. For $\alpha = 1.16$, the average value of a misreport steadily increases with n , as high-wager agents (who have high-value misreports) become more and more frequent, while low-value misreports become less frequent. With uniform wagers, the value of a misreport quickly decreases with n . With only 10 agents, a misreporting agent may be able to affect the price quite significantly, however with increasing n , misreports will consist of only being able to make small adjustments to the security price. For $\alpha = 3$, the value of a misreport remains approximately constant as n increases, suggesting some combination of the two previous effects.

Figure 1c shows how the average wager of agents with a profitable misreport changes with n . For uniform wagers this line is flat, since all agents have wager $w_i = 1$. The other two wager distributions display increasing wagers, which is again explained by increasing frequency of high-wager agents (with this frequency increasing faster for $\alpha = 1.16$ than for $\alpha = 3$), and decreasing frequency of low-wager agents that are actually able to profitably misreport.

Finally, in Figure 1d we plot the average distance between a profitable misreport and an agent's true belief. In contrast with the other statistics that we consider, this one is actually relatively flat as n increases (with the exception of a significant drop from $n = 10$ to $n = 20$). This tells us that even for small numbers of forecasters, misreporting is limited to agents with beliefs fairly close to the price π and does not significantly affect the equilibrium price. Note that the change in equilibrium price due to a misreport is upper bounded by the misreport distance, by Theorem 3.

We note that we have considered an omniscient setting where manipulating agents have precise knowledge of the reports of other agents. In practice, of course, the manipulating agent has uncertainty about her opponents. A misreport is risky, involving some possibility of being forced to buy securities at a price favorable to her misreport but not her true belief. High-wager agents have the most opportunities to misreport but also the most to lose if they miscalculate.

6 Conclusion

We have provided an axiomatic justification of the parimutuel consensus mechanism. While no wagering mechanism can satisfy anonymity, individual rationality, budget balance, Pareto optimality, sybilproofness, envy-freeness and incentive compatibility, we show that the PCM comes very close in that it satisfies all of the first six properties, and a relaxation of the seventh: incentive compatibility in the large. Subject to a mild condition on the reports, the PCM is the only wagering mechanism that satisfies all six properties. Via comprehensive simulations based on real contest data, we have shown that on large instances, opportunities to profitably manipulate are extremely rare. Even on small instances, the vast majority of agents cannot gain from misreporting.

Acknowledgments

The authors thank Jennifer Wortman Vaughan for helpful discussions. This work was completed in part while Freeman was an intern at Microsoft Research. He also thanks a Facebook Ph.D. Fellowship for support.

References

- [Ali, 1977] Mukhtar M. Ali. Probability and utility estimates for racetrack bettors. *Journal of Political Economy*, 85(4):803–816, 1977.
- [Arrow *et al.*, 2008] Kenneth J. Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz. The promise of prediction markets. *Science*, 320(5878):877–878, 2008.
- [Azevedo and Budish, 2017] Eduardo M Azevedo and Eric Budish. Strategy-proofness in the large. Technical report, National Bureau of Economic Research, 2017.
- [Beygelzimer *et al.*, 2012] Alina Beygelzimer, John Langford, and David M. Pennock. Learning performance of prediction markets with Kelly bettors. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1317–1318, 2012.
- [Chen and Pennock, 2007] Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 349–358, 2007.
- [Chen *et al.*, 2014] Yiling Chen, Nikhil R. Devanur, David M. Pennock, and Jennifer Wortman Vaughan. Removing arbitrage from wagering mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.
- [Cover and Thomas, 2006] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory, 2nd Edition*. Wiley, New York, 2006.
- [Ehlers *et al.*, 2004] Lars Ehlers, Hans Peters, and Ton Storcken. Threshold strategy-proofness: On manipulability in large voting problems. *Games and Economic Behavior*, 49(1):103–116, 2004.
- [Eisenberg and Gale, 1959] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The parimutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [Foley, 1967] D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- [Freeman *et al.*, 2017] Rupert Freeman, David M Pennock, and Jennifer Wortman Vaughan. The double clinching auction for wagering. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 43–60, 2017.
- [Galebach, 2004] Brian Galebach. Probabilitiesports. <http://probabilitiesports.com/>, 2004.
- [Hanson, 2003] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- [Lambert *et al.*, 2008] Nicolas S. Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, 2008.
- [Lambert *et al.*, 2015] Nicolas S. Lambert, John Langford, Jennifer Wortman Vaughan, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory*, 156:389–416, 2015.
- [Manski, 2006] Charles F. Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429, June 2006.
- [Plott *et al.*, 2003] Charles R Plott, Jorgen Wit, and Winston C Yang. Parimutuel betting markets as information aggregation devices: Experimental results. *Economic Theory*, 22(2):311–351, 2003.
- [Roberts and Postlewaite, 1976] Donald John Roberts and Andrew Postlewaite. The incentives for price-taking behavior in large exchange economies. *Econometrica: Journal of the Econometric Society*, pages 115–127, 1976.
- [Snyder, 1978] Wayne W. Snyder. Horse racing: Testing the efficient markets model. *Journal of Finance*, 33(4):1109–1118, 1978.
- [Thaler and Ziemba, 1988] Richard H. Thaler and William T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- [Wolfers and Zitzewitz, 2006] Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. Technical Report Working Paper 12200, National Bureau of Economic Research, 2006.