An Axiomatic View of the Parimutuel Consensus Wagering Mechanism

Extended Abstract

Rupert Freeman
Duke University
Durham, NC, USA
rupert@cs.duke.edu

David M. Pennock
Microsoft Research
New York, NY, USA
dpennock@microsoft.com

ABSTRACT

We consider an axiomatic view of the *Parimutuel Consensus Mechanism* defined by Eisenberg and Gale [6]. 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, 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.

ACM Reference Format:
Rupert Freeman and David M. Pennock. 2018. An Axiomatic View of the Parimutuel Consensus Wagering Mechanism. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018*, IFAAMAS, 3 pages.

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. 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 [6] analyzed the equilibrium of parimutuel betting, defining the *parimutuel consensus mechanism (PCM)*. The PCM is equivalent to parimutuel betting with 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 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 in question. Thus, facilitating wagering can serve as a source of freely provided information for a principal seeking to forecast a future event.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

Eisenberg and Gale discuss one undesirable feature of the equilibrium: it produces odds that sometimes ignore some agents. Manski [11] 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*—an agent's 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 [1, 13, 14], like many prediction market mechanisms do [2]. Plott et al. [12] 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 [4, 15]. In particular, an agent with logarithmic utility does best by betting an amount on each outcome proportional to her probability [5].

We examine another plausible reason why the PCM continues to enjoy usage: the mechanism satisfies a number of desirable axioms for wagering systems. Following the approach of Lambert et al. [9, 10], Freeman et al. [7] observed that it satisfies Pareto optimality (the mechanism exhausts all mutually beneficial trades), budget balance (the mechanism does not gain or lose money), and individual rationality (agents do not lose money in expectation). We further prove that the PCM satisfies sybilproofness, which ensures that one agent signing up as two, or controlling a shill account, cannot improve her lot.

Unfortunately, even the first three properties are mathematically incompatible with incentive compatibility [7]. Yet we can show that the PCM is near incentive compatible. 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 four natural and desirable properties of wagering mechanisms and it comes close both theoretically and empirically to obtaining a crucial fifth: incentive compatibility.

2 OUR CONTRIBUTIONS

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 regarding 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 $w_i \geq 0$ to the mechanism, where $\hat{p}_i$ captures her belief and w_i her budget. The mechanism outputs an allocation of Arrow-Debreu securities of the following form: A yes security pays off \$1 if $X = 1$ and \$0 otherwise, and a no security pays off \$0 if $X = 1$ and \$1 otherwise. Denote by y_i (resp. n_i) the number of yes (resp. no) securities allocated to agent i, and by σ_i the cost paid by i for those securities.

The *Parimutuel Consensus Mechanism (PCM)* [6] 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 probability less than π is allocated no securities at a price of $1 - \pi$ per security, and an agent with probability 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 possibility of needing agents with $\hat{p}_i = \pi$ to 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 [6] show as their main contribution that such a price is both unique and guaranteed to exist. The output of the PCM is defined by

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

and

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

It is known that the PCM satisfies Pareto optimality, budget balance, and individual rationality [7].[1] We show that it also satisfies sybilproofness.

[1]For formal definitions of these properties, as well as missing proofs and additional results, see the full version of the paper, which is available on the authors' websites.

	% Agents With Profitable Misreports	Average Profit
Pareto(α = 1.16)	0.07	1.55
Pareto(α = 3)	< 0.01	0.03
Uniform	0	n/a

Table 1: **Profitable misreports under Pareto and uniform wagers.**

PROPOSITION 2.1. *The parimutuel consensus mechanism satisfies sybilproofness.*

Our main theoretical result is that the PCM satisfies *incentive-compatibility in the large (IC-L)* [3]. Full incentive compatibility requires that truthful reporting is optimal for *every* realization of other agents' reports, but IC-L relaxes this condition in two ways. It requires 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 in the (ex-post) realization.

THEOREM 2.2. *The parimutuel consensus mechanism satisfies incentive compatibility in the large.*

To support Theorem 2.2, we test the PCM on a data set consisting of probability reports gathered from an online prediction contest called ProbabilitySports [8].[2] The data set consists of 1643 sets of probabilistic predictions, with each set containing between 64 and 1574 reports. Participants provided predictions but not wagers, so we generated wagers from three different Pareto distributions, which approximately model the distribution of wealth in a population. The $\alpha = 1.16$ distribution is classically viewed as a realistic distribution of wealth, and is described by the "80/20" rule: 20% of the population has 80% of the wealth. The $\alpha = 3$ distribution produces a more equal distribution of wagers, and the uniform distribution (which is the limit of the Pareto distribution as $\alpha \to \infty$) models situations where all wagers are equal, or where participants do not get to choose a wager. For every set of wagers generated, wagers are scaled so that the average wager is exactly 1.

Table 1 shows the percentage of agents that have a profitable misreport available, averaged across 50 random sets of wagers for each 1643 sets of predictions, and the average expected profit for those agents that have a profitable misreport. Strikingly, very few agents have an opportunity to profit – for uniform wagers, we did not find a single profitable misreport. The Pareto($\alpha = 1.16$) distribution has the most opportunities for misreports, and the highest profit misreports, because it has a higher fraction of high-wager agents. These agents have the most ability to affect the security price π in a favorable manner by reporting strategically.

The full version of the paper contains a more complete analysis. We find that even for smaller instances, with $10 \leq |\mathcal{N}| \leq 50$, profitable misreports are relatively rare, and their number decreases sharply as more agents are added. This suggests that the PCM may be a suitable choice of wagering mechanism in many real-life applications.

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.

[2]We thank Brian Galebach for providing us with this data.

REFERENCES

[1] Mukhtar M. Ali. 1977. Probability and utility estimates for racetrack bettors. *Journal of Political Economy* 85, 4 (1977), 803–816.

[2] 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. 2008. The promise of prediction markets. *Science* 320, 5878 (2008), 877–878.

[3] Eduardo M Azevedo and Eric Budish. 2017. *Strategy-Proofness in the Large*. Technical Report. National Bureau of Economic Research.

[4] Alina Beygelzimer, John Langford, and David M. Pennock. 2012. Learning Performance of Prediction Markets with Kelly Bettors. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1317–1318.

[5] Thomas M. Cover and Joy A. Thomas. 2006. *Elements of Information Theory, 2nd Edition*. Wiley, New York.

[6] Edmund Eisenberg and David Gale. 1959. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.

[7] Rupert Freeman, David M Pennock, and Jennifer Wortman Vaughan. 2017. The Double Clinching Auction for Wagering. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 43–60.

[8] Brian Galebach. 2004. ProbabilitySports. http://probabilitysports.com/. (2004).

[9] Nicolas S. Lambert, John Langford, Jennifer Wortman Vaughan, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. 2015. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory* 156 (2015), 389–416.

[10] Nicolas S. Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. 2008. Self-Financed Wagering Mechanisms for Forecasting. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*.

[11] Charles F. Manski. 2006. Interpreting the Predictions of Prediction Markets. *Economics Letters* 91, 3 (June 2006), 425–429.

[12] Charles R Plott, Jorgen Wit, and Winston C Yang. 2003. Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results. *Economic Theory* 22, 2 (2003), 311–351.

[13] Wayne W. Snyder. 1978. Horse Racing: Testing the Efficient Markets Model. *Journal of Finance* 33, 4 (1978), 1109–1118.

[14] Richard H. Thaler and William T. Ziemba. 1988. Anomalies: Parimutuel Betting Markets: Racetracks and lotteries. *Journal of Economic Perspectives* 2, 2 (1988), 161–174.

[15] Justin Wolfers and Eric Zitzewitz. 2006. *Interpreting Prediction Market Prices as Probabilities*. Technical Report Working Paper 12200. National Bureau of Economic Research.