

A Dynamic Pari-Mutuel Market for Hedging, Wagering, and Information Aggregation

David M. Pennock
Yahoo! Research Labs
74 N. Pasadena Ave, 3rd Floor
Pasadena, CA 91103 USA
pennockd@yahoo-inc.com

ABSTRACT

I develop a new mechanism for risk allocation and information speculation called a *dynamic pari-mutuel market* (DPM). A DPM acts as hybrid between a pari-mutuel market and a continuous double auction (CDA), inheriting some of the advantages of both. Like a pari-mutuel market, a DPM offers infinite buy-in liquidity and zero risk for the market institution; like a CDA, a DPM can continuously react to new information, dynamically incorporate information into prices, and allow traders to lock in gains or limit losses by selling prior to event resolution. The trader interface can be designed to mimic the familiar double auction format with bid-ask queues, though with an addition variable called the payoff per share. The DPM price function can be viewed as an automated market maker always offering to sell at some price, and moving the price appropriately according to demand. Since the mechanism is pari-mutuel (i.e., redistributive), it is guaranteed to pay out exactly the amount of money taken in. I explore a number of variations on the basic DPM, analyzing the properties of each, and solving in closed form for their respective price functions.

Categories and Subject Descriptors

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

General Terms

Algorithms, Design, Economics, Theory.

Keywords

Dynamic pari-mutuel market, continuous double auction, automated market maker, compound securities markets, combinatorial betting, risk allocation, information aggregation, trading, hedging, speculating, betting, wagering, gambling.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'04, May 17–20, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-711-0/04/0005 ...\$5.00.

1. INTRODUCTION

A wide variety of financial and wagering mechanisms have been developed to support hedging (i.e., insuring) against exposure to uncertain events and/or speculative trading on uncertain events. The dominant mechanism used in financial circles is the continuous double auction (CDA), or in some cases the *CDA with market maker* (CDAwMM). The primary mechanism used for sports wagering is a *bookie* or *bookmaker*, who essentially acts exactly as a market maker. Horse racing and jai alai wagering traditionally employ the *pari-mutuel* mechanism. Though there is no formal or logical separation between financial trading and wagering, the two endeavors are socially considered distinct. Recently, there has been a move to employ CDAs or CDAwMMs for all types of wagering, including on sports, horse racing, political events, world news, and many other uncertain events, and a simultaneous and opposite trend to use bookie systems for betting on financial markets. These trends highlight the interchangeable nature of the mechanisms and further blur the line between investing and betting. Some companies at the forefront of these movements are growing exponentially, with some industry observers declaring the onset of a revolution in the wagering business.¹

Each mechanism has pros and cons for the market institution and the participating traders. A CDA only matches willing traders, and so poses no risk whatsoever for the market institution. But a CDA can suffer from illiquidity in the form huge bid-ask spreads or even empty bid-ask queues if trading is light and thus markets are thin. A successful CDA must overcome a chicken-and-egg problem: traders are attracted to liquid markets, but liquid markets require a large number of traders. A CDAwMM and the similar bookie mechanism have built-in liquidity, but at a cost: the market maker itself, usually affiliated with the market institution, is exposed to significant risk of large monetary losses. Both the CDA and CDAwMM offer incentives for traders to leverage information continuously as soon as that information becomes available. As a result, prices are known to capture the current state of information exceptionally well.

Pari-mutuel markets effectively have infinite liquidity: anyone can place a bet on any outcome at any time, without the need for a matching offer from another bettor or a market maker. Pari-mutuel markets also involve no risk for the market institution, since they only redistribute money from losing wagers to winning wagers. However, pari-mutuel mar-

¹<http://www.wired.com/news/ebiz/0,1272,61051,00.html>

kets are not suitable for situations where information arrives over time, since there is a strong disincentive for placing bets until either (1) all information is revealed, or (2) the market is about to close. For this reason, pari-mutuel “prices” prior to the market’s close cannot be considered a reflection of current information. Pari-mutuel market participants cannot “buy low and sell high”: they cannot cash out gains (or limit losses) before the event outcome is revealed. Because the process whereby information arrives continuously over time is the rule rather than the exception, the applicability of the standard pari-mutuel mechanism is questionable in a large number of settings.

In this paper, I develop a new mechanism suitable for hedging, speculating, and wagering, called a *dynamic pari-mutuel market* (DPM). A DPM can be thought of as a hybrid between a pari-mutuel market and a CDA. A DPM is indeed pari-mutuel in nature, meaning that it acts only to redistribute money from some traders to others, and so exposes the market institution to no volatility (no risk). A constant, pre-determined subsidy is required to start the market. The subsidy can in principle be arbitrarily small and might conceivably come from traders (via antes or transaction fees) rather than the market institution, though a nontrivial outside subsidy may actually encourage trading and information aggregation. A DPM has the infinite liquidity of a pari-mutuel market: traders can always purchase shares in any outcome at any time, at some price automatically set by the market institution. A DPM is also able to react to and incorporate information arriving over time, like a CDA. The market institution changes the price for particular outcomes based on the current state of wagering. If a particular outcome receives a relatively large number of wagers, its price increases; if an outcome receives relatively few wagers, its price decreases. Prices are computed automatically using a *price function*, which can differ depending on what properties are desired. The price function determines the instantaneous price per share for an infinitesimal quantity of shares; the total cost for purchasing n shares is computed as the integral of the price function from 0 to n . The complexity of the price function can be hidden from traders by communicating only the ask prices for various lots of shares (e.g., lots of 100 shares), as is common practice in CDAs and CDawMMs. DPM prices do reflect current information, and traders can cash out in an aftermarket to lock in gains or limit losses before the event outcome is revealed. While there is always a market maker willing to accept buy orders, there is *not* a market maker accepting sell orders, and thus no guaranteed liquidity for selling: instead, selling is accomplished via a standard CDA mechanism. Traders can always “hedge-sell” by purchasing the opposite outcome than they already own.

2. BACKGROUND AND RELATED WORK

2.1 Pari-mutuel markets

Pari-mutuel markets are common at horse races [1, 22, 24, 25, 26], dog races, and jai alai games. In a pari-mutuel market people place wagers on which of two or more mutually exclusive and exhaustive outcomes will occur at some time in the future. After the true outcome becomes known, all of the money that is lost by those who bet on the incorrect outcome is redistributed to those who bet on the correct outcome, in direct proportion to the amount they wagered.

More formally, if there are k mutually exclusive and exhaustive outcomes (e.g., k horses, exactly one of which will win), and M_1, M_2, \dots, M_k dollars are bet on each outcome, and outcome i occurs, then everyone who bet on an outcome $j \neq i$ loses their wager, while everyone who bet on outcome i receives $\sum_{j=1}^k M_j/M_i$ dollars for every dollar they wagered. That is, every dollar wagered on i receives an equal share of all money wagered. An equivalent way to think about the redistribution rule is that every dollar wagered on i is refunded, then receives an equal share of all remaining money bet on the losing outcomes, or $\sum_{j \neq i} M_j/M_i$ dollars.

In practice, the market institution (e.g., the racetrack) first takes a certain percent of the total amount wagered, usually about 20% in the United States, then redistributes whatever money remains to the winners in proportion to their amount bet.

Consider a simple example with two outcomes, A and B . The outcomes are mutually exclusive and exhaustive, meaning that $\Pr(A \wedge B) = 0$ and $\Pr(A) + \Pr(B) = 1$. Suppose \$800 is bet on A and \$200 on B . Now suppose that A occurs (e.g., horse A wins the race). People who wagered on B lose their money, or \$200 in total. People who wagered on A win and each receives a proportional share of the total \$1000 wagered (ignoring fees). Specifically, each \$1 wager on A entitles its owner a $1/800$ share of the \$1000, or \$1.25.

Every dollar bet in a pari-mutuel market has an equal payoff, regardless of when the wager was placed or how much money was invested in the various outcomes at the time the wager was placed. The only state that matters is the final state: the final amounts wagered on all the outcomes when the market closes, and the identity of the correct outcome. As a result, there is a disincentive to place a wager early if there is any chance that new information might become available. Moreover, there are no guarantees about the payoff rate of a particular bet, except that it will be nonnegative if the correct outcome is chosen. Payoff rates can fluctuate arbitrarily until the market closes. So a second reason not to bet early is to wait to get a better sense of the final payout rates. This is in contrast to CDAs and CDawMMs, like the stock market, where incentives exist to invest as soon as new information is revealed.

Pari-mutuel bettors may be allowed to switch their chosen outcome, or even cancel their bet, prior to the market’s close. However, they cannot cash out of the market early, to either lock in gains or limit losses, if new information favors one outcome over another, as is possible in a CDA or a CDawMM. If bettors can cancel or change their bets, then an aftermarket to sell existing wagers is not sensible: every dollar wagered is worth exactly \$1 up until the market’s close—no one would buy at greater than \$1 and no one would sell at less than \$1. Pari-mutuel bettors must wait until the outcome is revealed to realize any profit or loss.

Unlike a CDA, in a pari-mutuel market, anyone can place a wager of any amount at any time—there is in a sense infinite liquidity for buying. A CDawMM also has built-in liquidity, but at the cost of significant risk for the market maker. In a pari-mutuel market, since money is only redistributed among bettors, the market institution itself has no risk. The main drawback of a pari-mutuel market is that it is useful only for capturing the value of an uncertain asset at some instant in time. It is ill-suited for situations where information arrives over time, continuously updating the estimated value of the asset—situations common in al-

most all trading and wagering scenarios. There is no notion of “buying low and selling high”, as occurs in a CDA, where buying when few others are buying (and the price is low) is rewarded more than buying when many others are buying (and the price is high). Perhaps for this reason, in most dynamic environments, financial mechanisms like the CDA that can react in real-time to changing information are more typically employed to facilitate speculating and hedging.

Since a pari-mutuel market can estimate the value of an asset at a single instant in time, a *repeated pari-mutuel market*, where distinct pari-mutuel markets are run at consecutive intervals, could in principle capture changing information dynamics. But running multiple consecutive markets would likely thin out trading in each individual market. Also, in each individual pari-mutuel market, the incentives would still be to wait to bet until just before the ending time of that particular market. This last problem might be mitigated by instituting a random stopping rule for each individual pari-mutuel market.

In laboratory experiments, pari-mutuel markets have shown a remarkable ability to aggregate and disseminate information dispersed among traders, at least for a single snapshot in time [17]. A similar ability has been recognized at real racetracks [1, 22, 24, 25, 26].

2.2 Financial markets

In the financial world, wagering on the outcomes of uncertain future propositions is also common. The typical market mechanism used is the continuous double auction (CDA). The term *securities market* in economics and finance generically encompasses a number of markets where speculating on uncertain events is possible. Examples include stock markets like NASDAQ, options markets like the CBOE [13], futures markets like the CME [21], other derivatives markets, insurance markets, political stock markets [6, 7], idea futures markets [12], decision markets [10] and even market games [3, 15, 16]. Securities markets generally have an economic and social value beyond facilitating speculation or wagering: they allow traders to *hedge risk*, or to insure against undesirable outcomes. So if a particular outcome has disutility for a trader, he or she can mitigate the risk by wagering *for* the outcome, to arrange for compensation in case the outcome occurs. In this sense, buying automobile insurance is effectively a bet that an accident or other covered event *will* occur. Similarly, buying a put option, which is useful as a hedge for a stockholder, is a bet that the underlying stock will go down. In practice, agents engage in a mixture of hedging and speculating, and there is no clear dividing line between the two [14]. Like pari-mutuel markets, often prices in financial markets are excellent information aggregators, yielding very accurate forecasts of future events [5, 18, 19].

A CDA constantly matches orders to buy an asset with orders to sell. If at any time one party is willing to buy one unit of the asset at a bid price of p_{bid} , while another party is willing to sell one unit of the asset at an ask price of p_{ask} , and p_{bid} is greater than or equal to p_{ask} , then the two parties transact (at some price between p_{bid} and p_{ask}). If the highest bid price is less than the lowest ask price, then no transactions occur. In a CDA, the bid and ask prices rapidly change as new information arrives and traders reassess the value of the asset. Since the auctioneer only matches willing bidders, the auctioneer takes on no risk. However, buyers

can only buy as many shares as sellers are willing to sell; for any transaction to occur, there must be a counterparty on the other side willing to accept the trade.

As a result, when few traders participate in a CDA, it may become illiquid, meaning that not much trading activity occurs. The spread between the highest bid price and the lowest ask price may be very large, or one or both queues may be completely empty, discouraging trading.² One way to induce liquidity is to provide a market maker who is willing to accept a large number of buy and sell orders at particular prices. We call this mechanism a *CDA with market maker* (CDAwMM).³ Conceptually, the market maker is just like any other trader, but typically is willing to accept a much larger volume of trades. The market maker may be a person, or may be an automated algorithm. Adding a market maker to the system increases liquidity, but exposes the market maker to risk. Now, instead of only matching trades, the system actually takes on risk of its own, and depending on what happens in the future, may lose considerable amounts of money.

2.3 Wagering markets

The typical Las Vegas bookmaker or oddsmaker functions much like a market maker in a CDA. In this case, the market institution (the *book* or *house*) sets the odds,⁴ initially according to expert opinion, and later in response to the relative level of betting on the various outcomes. Unlike in a pari-mutuel environment, whenever a wager is placed with a bookmaker, the odds or terms for that bet are fixed at the time of the bet. The bookmaker profits by offering different odds for the two sides of the bet, essentially defining a bid-ask spread. While odds may change in response to changing information, any bets made at previously set odds remain in effect according to the odds at the time of the bet; this is precisely in analogy to a CDAwMM. One difference between a bookmaker and a market maker is that the former usually operates in a “take it or leave it mode”: bettors cannot place their own limit orders on a common queue, they can in effect only place market orders at prices defined by the bookmaker. Still, the bookmaker certainly reacts to bettor demand. Like a market maker, the bookmaker exposes itself to significant risk. Sports betting markets have also been shown to provide high quality aggregate forecasts [4, 9, 23].

2.4 Market scoring rule

Hanson’s [11] *market scoring rule* (MSR) is a new mechanism for hedging and speculating that shares some properties in common with a DPM. Like a DPM, an MSR can be conceptualized as an automated market maker always willing to accept a trade on any event at some price. An MSR requires a patron to subsidize the market. The patron’s final loss is variable, and thus technically implies a degree of risk, though the maximum loss is bounded. An MSR maintains a probability distribution over all events. At any time any

²Thin markets do occur often in practice, and can be seen in a variety of the less popular markets available on <http://TradeSports.com>, or in some financial options markets, for example.

³A very clear example of a CDAwMM is the “interactive” betting market on <http://WSEX.com>.

⁴Or, alternatively, the bookmaker sets the game *line* in order to provide even-money odds.

trader who believes the probabilities are wrong can change any part of the distribution by accepting a lottery ticket that pays off according to a scoring rule (e.g., the logarithmic scoring rule) [27], as long as that trader also agrees to pay off the most recent person to change the distribution. In the limit of a single trader, the mechanism behaves like a scoring rule, suitable for polling a single agent for its probability distribution. In the limit of many traders, it produces a combined estimate. Since the market essentially always has a complete set of posted prices for all possible outcomes, the mechanism avoids the problem of thin markets or illiquidity. An MSR is not pari-mutuel in nature, as the patron in general injects a variable amount of money into the system. An MSR provides a two-sided automated market maker, while a DPM provides a one-sided automated market maker. In an MSR, the vector of payoffs across outcomes is fixed at the time of the trade, while in a DPM, the vector of payoffs across outcomes depends both on the state of wagering at the time of the trade and the state of wagering at the market's close. While the mechanisms are quite different—and so trader acceptance and incentives may strongly differ—the properties and motivations of DPMs and MSRs are quite similar.

Hanson shows how MSRs are especially well suited for allowing bets on a combinatorial number of outcomes. The patron's payment for subsidizing trading on all 2^n possible combinations of n events is no larger than the sum of subsidizing the n event marginals independently. The mechanism was planned for use in the *Policy Analysis Market* (PAM), a futures market in Middle East related outcomes and funded by DARPA [20], until a media firestorm killed the project.⁵ As of this writing, the founders of PAM were considering reopening under private control.⁶

3. A DYNAMIC PARI-MUTUEL MARKET

3.1 High-level description

In contrast to a standard pari-mutuel market, where each dollar always buys an equal share of the payoff, in a DPM each dollar buys a variable share in the payoff depending on the state of wagering at the time of purchase. So a wager on A at a time when most others are wagering on B offers a greater possible profit than a wager on A when most others are also wagering on A .

A natural way to communicate the changing payoff of a bet is to say that, at any given time, a certain amount of money will buy a certain number of *shares* in one outcome the other. Purchasing a share entitles its owner to an equal stake in the winning pot should the chosen outcome occur. The payoff is variable, because when few people are betting on an outcome, shares will generally be cheaper than at a time when many people are betting that outcome. There is no pre-determined limit on the number of shares: new shares can be continually generated as trading proceeds.

For simplicity, all analyses in this paper consider the binary outcome case; generalizing to multiple discrete outcomes should be straightforward. Denote the two outcomes A and B . The outcomes are mutually exclusive and ex-

haustive. Denote the instantaneous price per share of A as p_1 and the price per share of B as p_2 . Denote the payoffs per share as \mathcal{P}_1 and \mathcal{P}_2 , respectively. These four numbers, $p_1, p_2, \mathcal{P}_1, \mathcal{P}_2$ are the key numbers that traders must track and understand. Note that the price is set at the time of the wager; the payoff per share is finalized only after the event outcome is revealed.

At any time, a trader can purchase an infinitesimal quantity of shares of A at price p_1 (and similarly for B). However, since the price changes continuously as shares are purchased, the cost of buying n shares is computed as the integral of a *price function* from 0 to n . The use of continuous functions and integrals can be hidden from traders by aggregating the automated market maker's sell orders into discrete lots of, say, 100 shares each. These ask orders can be automatically entered into the system by the market institution, so that traders interact with what looks like a more familiar CDA; we examine this interface issue in more detail below in Section 4.2.

For our analysis, we introduce the following additional notation. Denote M_1 as the total amount of money wagered on A , M_2 as the total amount of money wagered on B , $T = M_1 + M_2$ as the total amount of money wagered on both sides, N_1 as the total number of shares purchased of A , and N_2 as the total number of shares purchased of B .

There are many ways to formulate the price function. Several natural price functions are outlined below; each is motivated as the unique solution to a particular constraint on price dynamics.

3.2 Advantages and disadvantages

To my knowledge, a DPM is the only known mechanism for hedging and speculating that exhibits all three of the following properties: (1) guaranteed liquidity, (2) no risk for the market institution, and (3) continuous incorporation of information. A standard pari-mutuel fails (3). A CDA fails (1). A CDawMM, the bookmaker mechanism, and an MSR all fail (2). Even though technically an MSR exposes its patron to risk (i.e., a variable future payoff), the patron's maximum loss is bounded, so the distinction between a DPM and an MSR in terms of these three properties is more technical than practical.

DPM traders can cash out of the market early, just like stock market traders, to lock in a profit or limit a loss, an action that is simply not possible in a standard pari-mutuel.

A DPM also has some drawbacks. The payoff for a wager depends both on the price at the time of the trade, and on the final payoff per share at the market's close. This contrasts with the CDA variants, where the payoff vector across possible future outcomes is fixed at the time of the trade. So a trader's strategic optimization problem is complicated by the need to predict the final values of \mathcal{P}_1 and \mathcal{P}_2 . If \mathcal{P} changes according to a random walk, then traders can take the current \mathcal{P} as an unbiased estimate of the final \mathcal{P} , greatly decreasing the complexity of their optimization. If \mathcal{P} does not change according to a random walk, the mechanism still has utility as a mechanism for hedging and speculating, though optimization may be difficult, and determining a measure of the market's aggregate opinion of the probabilities of A and B may be difficult. We discuss the implications of random walk behavior further below in Section 4.1 in the discussion surrounding Assumption 3.

A second drawback of a DPM is its one-sided nature.

⁵See <http://hanson.gmu.edu/policyanalysismarket.html> for more information, or <http://dpennock.com/pam.html> for commentary.

⁶<http://www.policyanalysismarket.com/>

While an automated market maker always stands ready to accept buy orders, there is no corresponding market maker to accept sell orders. Traders must sell to each other using a standard CDA mechanism, for example by posting an ask order at a price at or below the market maker's current ask price. Traders can also always "hedge-sell" by purchasing shares in the opposite outcome from the market maker, thereby hedging their bet if not fully liquidating it.

3.3 Redistribution rule

In a standard pari-mutuel market, payoffs can be computed in either of two equivalent ways: (1) each winning \$1 wager receives a *refund* of the initial \$1 paid, plus an equal share of all losing wagers, or (2) each winning \$1 wager receives an equal share of *all* wagers, winning or losing. Because each dollar always earns an equal share of the payoff, the two formulations are precisely the same:

$$\$1 + \frac{M_{\text{lose}}}{M_{\text{win}}} = \frac{M_{\text{win}} + M_{\text{lose}}}{M_{\text{win}}}.$$

In a dynamic pari-mutuel market, because each dollar is not equally weighted, the two formulations are distinct, and lead to significantly different price functions and mechanisms, each with different potentially desirable properties. We consider each case in turn. The next section analyzes case (1), where only losing money is redistributed. Section 5 examines case (2), where all money is redistributed.

4. DPM I: LOSING MONEY REDISTRIBUTED

For the case where the initial payments on winning bets are refunded, and only losing money is redistributed, the respective payoffs per share are simply:

$$\begin{aligned} \mathcal{P}_1 &= \frac{M_2}{N_1} \\ \mathcal{P}_2 &= \frac{M_1}{N_2}. \end{aligned}$$

So, if A occurs, shareholders of A receive all of their initial payment back, plus \mathcal{P}_1 dollars per share owned, while shareholders of B lose all money wagered. Similarly, if B occurs, shareholders of B receive all of their initial payment back, plus \mathcal{P}_2 dollars per share owned, while shareholders of A lose all money wagered.

Without loss of generality, I will analyze the market from the perspective of A , deriving prices and payoffs for A only. The equations for B are symmetric.

The trader's per-share expected value for purchasing an infinitesimal quantity ϵ of shares of A is

$$\begin{aligned} \frac{E[\epsilon \text{ shares}]}{\epsilon} &= \Pr(A) \cdot E[\mathcal{P}_1|A] - (1 - \Pr(A)) \cdot p_1 \\ \frac{E[\epsilon \text{ shares}]}{\epsilon} &= \Pr(A) \cdot E\left[\frac{M_2}{N_1} \middle| A\right] - (1 - \Pr(A)) \cdot p_1 \end{aligned}$$

where ϵ is an infinitesimal quantity of shares of A , $\Pr(A)$ is the trader's belief in the probability of A , and p_1 is the instantaneous price per share of A for an infinitesimal quantity of shares. $E[\mathcal{P}_1|A]$ is the trader's expectation of the payoff per share of A after the market closes and *given that A occurs*. This is a subtle point. The value of \mathcal{P}_1 does not matter if B occurs, since in this case shares of A are worthless, and the *current* value of \mathcal{P}_1 does not necessarily matter

as this may change as trading continues. So, in order to determine the expected value of shares of A , the trader must estimate what he or she expects the payoff per share to be in the end (after the market closes) if A occurs.

If $E[\epsilon \text{ shares}]/\epsilon > 0$, a risk-neutral trader should purchase shares of A . How many shares? This depends on the price function determining p_1 . In general, p_1 increases as more shares are purchased. The risk-neutral trader should continue purchasing shares until $E[\epsilon \text{ shares}]/\epsilon = 0$. (A risk-averse trader will generally stop purchasing shares before driving $E[\epsilon \text{ shares}]/\epsilon$ all the way to zero.) Assuming risk-neutrality, the trader's optimization problem is to choose a number of shares $n \geq 0$ of A to purchase, in order to maximize

$$E[n \text{ shares}] = \Pr(A) \cdot n \cdot E[\mathcal{P}_1|A] - (1 - \Pr(A)) \cdot \int_0^n p_1(n) dn. \quad (1)$$

It's easy to see that the same value of n can be solved for by finding the number of shares required to drive $E[\epsilon \text{ shares}]/\epsilon$ to zero. That is, find $n \geq 0$ satisfying

$$0 = \Pr(A) \cdot E[\mathcal{P}_1|A] - (1 - \Pr(A)) \cdot p_1(n),$$

if such a n exists, otherwise $n = 0$.

4.1 Market probability

As traders who believe that $E[\epsilon \text{ shares of } A]/\epsilon > 0$ purchase shares of A and traders who believe that $E[\epsilon \text{ shares of } B]/\epsilon > 0$ purchase shares of B , the prices p_1 and p_2 change according to a price function, as prescribed below. The current prices in a sense reflect the market's opinion as a whole of the relative probabilities of A and B . Assuming an efficient marketplace, the market as a whole considers $E[\epsilon \text{ shares}]/\epsilon = 0$, since the mechanism is a zero sum game. For example, if market participants in aggregate felt that $E[\epsilon \text{ shares}]/\epsilon > 0$, then there would be net demand for A , driving up the price of A until $E[\epsilon \text{ shares}]/\epsilon = 0$. Define $\text{MPr}(A)$ to be the *market probability* of A , or the probability of A inferred by assuming that $E[\epsilon \text{ shares}]/\epsilon = 0$. We can consider $\text{MPr}(A)$ to be the aggregate probability of A as judged by the market as a whole. $\text{MPr}(A)$ is the solution to

$$0 = \text{MPr}(A) \cdot E[\mathcal{P}_1|A] - (1 - \text{MPr}(A)) \cdot p_1.$$

Solving we get

$$\text{MPr}(A) = \frac{p_1}{p_1 + E[\mathcal{P}_1|A]}. \quad (2)$$

At this point we make a critical assumption in order to greatly simplify the analysis; we assume that

$$E[\mathcal{P}_1|A] = \mathcal{P}_1. \quad (3)$$

That is, we assume that the current value for the payoff per share of A is the same as the expected final value of the payoff per share of A given that A occurs. This is certainly true for the last (infinitesimal) wager before the market closes. It's not obvious, however, that the assumption is true well before the market's close. Basically, we are assuming that the value of \mathcal{P}_1 moves according to an unbiased random walk: the current value of \mathcal{P}_1 is the best expectation of its future value. I conjecture that there are reasonable market efficiency conditions under which assumption (3) is true, though I have not been able to prove that it arises naturally from rational trading. We examine scenarios below in which

assumption (3) seems especially plausible. Nonetheless, the assumption effects our analysis only. Regardless of whether (3) is true, each price function derived below implies a well-defined zero-sum game in which traders can play. If traders can assume that (3) is true, then their optimization problem (1) is greatly simplified; however, optimizing (1) does not depend on the assumption, and traders can still optimize by strategically projecting the final expected payoff in whatever complicated way they desire. So, the utility of DPM for hedging and speculating does not necessarily hinge on the truth of assumption (3). On the other hand, the ability to easily infer an aggregate market consensus probability from market prices does depend on (3).

4.2 Price functions

A variety of price functions seem reasonable, each exhibiting various properties, and implying differing market probabilities.

4.2.1 Price function I: Price of A equals payoff of B

One natural price function to consider is to set the price per share of A equal to the payoff per share of B, and set the price per share of B equal to the payoff per share of A. That is,

$$\begin{aligned} p_1 &= \mathcal{P}_2 \\ p_2 &= \mathcal{P}_1. \end{aligned} \quad (4)$$

Enforcing this relationship reduces the dimensionality of the system from four to two, simplifying the interface: traders need only track two numbers instead of four. The relationship makes sense, since new information supporting A should encourage purchasing of shares A, driving up both the price of A and the payoff of B, and driving down the price of B and the payoff of A. In this setting, assumption (3) seems especially reasonable, since if an efficient market hypothesis leads prices to follow a random walk, than payoffs must also follow a random walk.

The constraints (4) lead to the following derivation of the market probability:

$$\begin{aligned} \text{MPr}(A)\mathcal{P}_1 &= \text{MPr}(B)p_1 \\ \text{MPr}(A)\mathcal{P}_1 &= \text{MPr}(B)\mathcal{P}_2 \\ \frac{\text{MPr}(A)}{\text{MPr}(B)} &= \frac{\mathcal{P}_2}{\mathcal{P}_1} \\ \frac{\text{MPr}(A)}{\text{MPr}(B)} &= \frac{\frac{M_1}{N_2}}{\frac{M_2}{N_1}} \\ \frac{\text{MPr}(A)}{\text{MPr}(B)} &= \frac{M_1 N_1}{M_2 N_2} \\ \text{MPr}(A) &= \frac{M_1 N_1}{M_1 N_1 + M_2 N_2} \end{aligned} \quad (5)$$

The constraints (4) specify the instantaneous relationship between payoff and price. From this, we can derive how prices change when (non-infinitesimal) shares are purchased. Let n be the number of shares purchased and let m be the amount of money spent purchasing n shares. Note that $p_1 = dm/dn$, the instantaneous price per share, and

$m = \int_0^n p_1(n)dn$. Substituting into equation (4), we get:

$$\begin{aligned} p_1 &= \mathcal{P}_2 \\ \frac{dm}{dn} &= \frac{M_1 + m}{N_2} \\ \frac{dm}{M_1 + m} &= \frac{dn}{N_2} \\ \int \frac{dm}{M_1 + m} &= \int \frac{dn}{N_2} \\ \ln(M_1 + m) &= \frac{n}{N_2} + C \\ m &= M_1 \left[e^{\frac{n}{N_2}} - 1 \right] \end{aligned} \quad (6)$$

Equation 6 gives the cost of purchasing n shares. The instantaneous price per share as a function of n is

$$p_1(n) = \frac{dm}{dn} = \frac{M_1}{N_2} e^{\frac{n}{N_2}}. \quad (7)$$

Note that $p_1(0) = M_1/N_2 = \mathcal{P}_2$ as required. The derivation of the price function $p_2(n)$ for B is analogous and the results are symmetric.

The notion of buying infinitesimal shares, or integrating costs over a continuous function, are probably foreign to most traders. A more standard interface can be implemented by discretizing the costs into round lots of shares, for example lots of 100 shares. Then ask orders of 100 shares each at the appropriate price can be automatically placed by the market institution. For example, the market institution can place an ask order for 100 shares at price $m(100)/100$, another ask order for 100 shares at price $(m(200) - m(100))/100$, a third ask for 100 shares at $(m(300) - m(200))/100$, etc. In this way, the market looks more familiar to traders, like a typical CDA with a number of ask orders at various prices automatically available. A trader buying less than 100 shares would pay a bit more than if the true cost were computed using (6), but the discretized interface would probably be more intuitive and transparent to the majority of traders.

The above equations assume that all money that comes in is eventually returned or redistributed. In other words, the mechanism is a zero sum game, and the market institution takes no portion of the money. This could be generalized so that the market institution always takes a certain amount, or a certain percent, or a certain amount per transaction, or a certain percent per transaction, before money is returned or redistributed.

Finally, note that the above price function is undefined when the amount bet or the number of shares are zero. So the system must begin with some positive amount on both sides, and some positive number of shares outstanding on both sides. These initial amounts can be arbitrarily small in principle, but the size of the initial subsidy may affect the incentives of traders to participate. Also, the smaller the initial amounts, the more each new dollar effects the prices. The initialization amounts could be funded as a subsidy from the market institution or a patron, which I'll call a *seed wager*, or from a portion of the fees charged, which I'll call an *ante wager*.

4.2.2 Price function II: Price of A proportional to money on A

A second price function can be derived by requiring the ratio of prices to be equal to the ratio of money wagered.

That is,

$$\frac{p_1}{p_2} = \frac{M_1}{M_2}. \quad (8)$$

In other words, the price of A is proportional to the amount of money wagered on A , and similarly for B . This seems like a particularly natural way to set the price, since the more money that is wagered on one side, the cheaper becomes a share on the other side, in exactly the same proportion.

Using Equation 8, along with (2) and (3), we can derive the implied market probability:

$$\begin{aligned} \frac{M_1}{M_2} &= \frac{p_1}{p_2} \\ &= \frac{\frac{\text{MPr}(A)}{\text{MPr}(B)} \cdot \frac{M_2}{N_1}}{\frac{\text{MPr}(B)}{\text{MPr}(A)} \cdot \frac{M_1}{N_2}} \\ &= \frac{(\text{MPr}(A))^2}{(\text{MPr}(B))^2} \cdot \frac{M_2 N_2}{M_1 N_1} \\ \frac{(\text{MPr}(A))^2}{(\text{MPr}(B))^2} &= \frac{(M_1)^2 N_1}{(M_2)^2 N_2} \\ \frac{\text{MPr}(A)}{\text{MPr}(B)} &= \frac{M_1 \sqrt{N_1}}{M_2 \sqrt{N_2}} \\ \text{MPr}(A) &= \frac{M_1 \sqrt{N_1}}{M_1 \sqrt{N_1} + M_2 \sqrt{N_2}} \end{aligned} \quad (9)$$

We can solve for the instantaneous price as follows:

$$\begin{aligned} p_1 &= \frac{\text{MPr}(A)}{\text{MPr}(B)} \cdot \mathcal{P}_1 \\ &= \frac{M_1 \sqrt{N_1}}{M_2 \sqrt{N_2}} \cdot \frac{M_2}{N_1} \\ &= \frac{M_1}{\sqrt{N_1 N_2}} \end{aligned} \quad (10)$$

Working from the above instantaneous price, we can derive the implied cost function m as a function of the number n of shares purchased as follows:

$$\begin{aligned} \frac{dm}{dn} &= \frac{M_1 + m}{\sqrt{N_1 + n} \sqrt{N_2}} \\ \int \frac{dm}{M_1 + m} &= \int \frac{dn}{\sqrt{N_1 + n} \sqrt{N_2}} \\ \ln(M_1 + m) &= \frac{2}{N_2} [(N_1 + n)N_2]^{\frac{1}{2}} + C \\ m &= M_1 \left[e^{2\sqrt{\frac{N_1+n}{N_2}} - 2\sqrt{\frac{N_1}{N_2}}} - 1 \right]. \end{aligned} \quad (11)$$

From this we get the price function:

$$p_1(n) = \frac{dm}{dn} = \frac{M_1}{\sqrt{(N_1 + n)N_2}} e^{2\sqrt{\frac{N_1+n}{N_2}} - 2\sqrt{\frac{N_1}{N_2}}}. \quad (12)$$

Note that, as required, $p_1(0) = M_1/\sqrt{N_1 N_2}$, and $p_1(0)/p_2(0) = M_1/M_2$. If one uses the above price function, then the market dynamics will be such that the ratio of the (instantaneous) prices of A and B always equals the ratio of the amounts wagered on A and B , which seems fairly natural.

Note that, as before, the mechanism can be modified to collect transaction fees of some kind. Also note that seed or ante wagers are required to initialize the system.

5. DPM II: ALL MONEY REDISTRIBUTED

Above we examined the policy of refunding winning wagers and redistributing only losing wagers. In this section we consider the second policy mentioned in Section 3.3: *all* money from all wagers are collected and redistributed to winning wagers.

For the case where all money is redistributed, the respective payoffs per share are:

$$\begin{aligned} \mathcal{P}_1 &= \frac{M_1 + M_2}{N_1} = \frac{T}{N_1} \\ \mathcal{P}_2 &= \frac{M_1 + M_2}{N_2} = \frac{T}{N_2}, \end{aligned}$$

where $T = M_1 + M_2$ is the total amount of money wagered on both sides. So, if A occurs, shareholders of A lose their initial price paid, but receive \mathcal{P}_1 dollars per share owned; shareholders of B simply lose all money wagered. Similarly, if B occurs, shareholders of B lose their initial price paid, but receive \mathcal{P}_2 dollars per share owned; shareholders of A lose all money wagered.

In this case, the trader's per-share expected value for purchasing an infinitesimal quantity ϵ of shares of A is

$$\frac{E[\epsilon \text{ shares}]}{\epsilon} = \text{Pr}(A) \cdot E[\mathcal{P}_1|A] - p_1. \quad (13)$$

A risk-neutral trader optimizes by choosing a number of shares $n \geq 0$ of A to purchase, in order to maximize

$$\begin{aligned} E[n \text{ shares}] &= \text{Pr}(A) \cdot n \cdot E[\mathcal{P}_1|A] - \int_0^n p_1(n) dn \\ &= \text{Pr}(A) \cdot n \cdot E[\mathcal{P}_1|A] - m \end{aligned} \quad (14)$$

The same value of n can be solved for by finding the number of shares required to drive $E[\epsilon \text{ shares}]/\epsilon$ to zero. That is, find $n \geq 0$ satisfying

$$0 = \text{Pr}(A) \cdot E[\mathcal{P}_1|A] - p_1(n),$$

if such a n exists, otherwise $n = 0$.

5.1 Market probability

In this case $\text{MPr}(A)$, the aggregate probability of A as judged by the market as a whole, is the solution to

$$0 = \text{MPr}(A) \cdot E[\mathcal{P}_1|A] - p_1.$$

Solving we get

$$\text{MPr}(A) = \frac{p_1}{E[\mathcal{P}_1|A]}. \quad (15)$$

As before, we make the simplifying assumption (3) that the expected final payoff per share equals the current payoff per share. The assumption is critical for our analysis, but may not be required for a practical implementation.

5.2 Price functions

For the case where all money is distributed, the constraints (4) that keep the price of A equal to the payoff of B , and vice versa, do not lead to the derivation of a coherent price function.

A reasonable price function can be derived from the constraint (8) employed in Section 4.2.2, where we require that the ratio of prices to be equal to the ratio of money wagered. That is, $p_1/p_2 = M_1/M_2$. In other words, the price of A is proportional to the amount of money wagered on A , and similarly for B .

Using Equations 3, 8, and 15 we can derive the implied market probability:

$$\begin{aligned}
\frac{M_1}{M_2} &= \frac{p_1}{p_2} \\
&= \frac{\text{MPr}(A)}{\text{MPr}(B)} \cdot \frac{T}{N_1} \cdot \frac{N_2}{T} \\
&= \frac{\text{MPr}(A)}{\text{MPr}(B)} \cdot \frac{N_2}{N_1} \\
\frac{\text{MPr}(A)}{\text{MPr}(B)} &= \frac{M_1 N_1}{M_2 N_2} \\
\text{MPr}(A) &= \frac{M_1 N_1}{M_1 N_1 + M_2 N_2} \quad (16)
\end{aligned}$$

Interestingly, this is the same market probability derived in Section 4.2.1 for the case of losing-money redistribution with the constraints that the price of A equal the payoff of B and vice versa.

The instantaneous price per share for an infinitesimal quantity of shares is:

$$\begin{aligned}
p_1 &= \frac{(M_1)^2 + M_1 M_2}{M_1 N_1 + M_2 N_2} \\
&= \frac{M_1 + M_2}{N_1 + \frac{M_2}{M_1} N_2}
\end{aligned}$$

Working from the above instantaneous price, we can derive the number of shares n that can be purchased for m dollars, as follows:

$$\begin{aligned}
\frac{dm}{dn} &= \frac{M_1 + M_2 + m}{N_1 + n + \frac{M_2}{M_1 + m} N_2} \\
\frac{dn}{dm} &= \frac{N_1 + n + \frac{M_2}{M_1 + m} N_2}{M_1 + M_2 + m} \quad (17) \\
&\dots \\
n &= \frac{m(N_1 - N_2)}{T} + \frac{N_2(T + m)}{M_2} \ln \left[\frac{T(M_1 + m)}{M_1(T + m)} \right].
\end{aligned}$$

Note that we solved for $n(m)$ rather than $m(n)$. I could not find a closed-form solution for $m(n)$, as was derived for the two other cases above. Still, $n(m)$ can be used to determine how many shares can be purchased for m dollars, and the inverse function can be approximated to any degree numerically. From $n(m)$ we can also compute the price function:

$$p_1(m) = \frac{dm}{dn} = \frac{(M_1 + m)M_2 T}{\text{denom}}, \quad (18)$$

where

$$\begin{aligned}
\text{denom} &= (M_1 + m)M_2 N_1 + (M_2 - m)M_2 N_2 \\
&\quad + T(M_1 + m)N_2 \ln \left[\frac{T(M_1 + m)}{M_1(T + m)} \right]
\end{aligned}$$

Note that, as required, $p_1(0)/p_2(0) = M_1/M_2$. If one uses the above price function, then the market dynamics will be such that the ratio of the (instantaneous) prices of A and B always equals the ratio of the amounts wagered on A and B .

This price function has another desirable property: it acts such that the expected value of wagering \$1 on A and simultaneously wagering \$1 on B equals zero, assuming (3). That is, $E[\$1 \text{ of } A + \$1 \text{ of } B] = 0$. The derivation is omitted.

5.3 Comparing DPM I and II

The main advantage of refunding winning wagers (DPM I) is that every bet on the winning outcome is guaranteed to at least break even. The main disadvantage of refunding winning wagers is that shares are not homogenous: each share of A , for example, is actually composed of two distinct parts: (1) the refund, or a lottery ticket that pays $\$p$ if A occurs, where p is the price paid per share, and (2) one share of the final payoff ($\$P_1$) if A occurs. This complicates the implementation of an aftermarket to cash out of the market early, which we will examine below in Section 7. When all money is redistributed (DPM II), shares are homogeneous: each share entitles its owner to an equal slice of the final payoff. Because shares are homogenous, the implementation of an aftermarket is straightforward, as we shall see in Section 7. On the other hand, because initial prices paid are not refunded for winning bets, there is a chance that, if prices swing wildly enough, a wager on the correct outcome might actually lose money. Traders must be aware that if they buy in at an excessively high price that later tumbles allowing many others to get in at a much lower price, they may lose money in the end regardless of the outcome. From informal experiments, I don't believe this eventuality would be common, but nonetheless it requires care in communicating to traders the possible risks. One potential fix would be for the market maker to keep track of when the price is going too low, endangering an investor on the correct outcome. At this point, the market maker could artificially stop lowering the price. Sell orders in the aftermarket might still come in below the market maker's price, but in this way the system could ensure that every wager on the correct outcome at least breaks even.

6. OTHER VARIATIONS

A simple ascending price function would set $p_1 = \alpha M_1$ and $p_2 = \alpha M_2$, where $\alpha > 0$. In this case, prices would only go up. For the case of all money being redistributed, this would eliminate the possibility of losing money on a wager on the correct outcome. Even though the market maker's price only rises, the going price may fall well below the market maker's price, as ask orders are placed in the aftermarket.

I have derived price functions for several other cases, using the same methodology above. Each price function may have its own desirable properties, but it's not clear which is best, or even that a single best method exists. Further analyses and, more importantly, empirical investigations are required to answer these questions.

7. AFTERMARKETS

A key advantage of DPM over a standard pari-mutuel market is the ability to cash out of the market before it closes, in order to take a profit or limit a loss. This is accomplished by allowing traders to place ask orders on the same queue as the market maker. So traders can sell the shares that they purchased at or below the price set by the market maker. Or traders can place a limit sell order at any price. Buyers will purchase any existing shares for sale at the lower prices first, before purchasing new shares from the market maker.

7.1 Aftermarket for DPM II

For the second main case explored above, where all money

is redistributed, allowing an aftermarket is simple. In fact, “aftermarket” may be a poor descriptor: buying and selling are both fully integrated into the same mechanism. Every share is worth precisely the same amount, so traders can simply place ask orders on the same queue as the market maker in order to sell their shares. New buyers will accept the lowest ask price, whether it comes from the market maker or another trader. In this way, traders can cash out early and walk away with their current profit or loss, assuming they can find a willing buyer.

7.2 Aftermarket for DPM I

When winning wagers are refunded and only losing wagers are redistributed, each share is potentially worth a different amount, depending on how much was paid for it, so it is not as simple a matter to set up an aftermarket. However, an aftermarket is still possible. In fact, much of the complexity can be hidden from traders, so it looks nearly as simple as placing a sell order on the queue.

In this case shares are not homogenous: each share of A is actually composed of two distinct parts: (1) the refund of $p \cdot 1_A$ dollars, and (2) the payoff of $\mathcal{P}_1 \cdot 1_A$ dollars, where p is the per-share price paid and 1_A is the indicator function equalling 1 if A occurs, and 0 otherwise. One can imagine running two separate aftermarkets where people can sell these two respective components. However, it is possible to automate the two aftermarkets, by automatically bundling them together in the correct ratio and selling them in the central DPM. In this way, traders can cash out by placing sell orders on the same queue as the DPM market maker, effectively hiding the complexity of explicitly having two separate aftermarkets. The bundling mechanism works as follows. Suppose the current price for 1 share of A is p_1 . A buyer agrees to purchase the share at p_1 . The buyer pays p_1 dollars and receives $p_1 \cdot 1_A + \mathcal{P}_1 \cdot 1_A$ dollars. If there is enough inventory in the aftermarkets, the buyer’s share is constructed by bundling together $p_1 \cdot 1_A$ from the first aftermarket, and $\mathcal{P}_1 \cdot 1_A$ from the second aftermarket. The seller in the first aftermarket receives $p_1 \text{MPr}(A)$ dollars, and the seller in the second aftermarket receives $p_1 \text{MPr}(B)$ dollars.

7.3 Pseudo aftermarket for DPM I

There is an alternative “pseudo aftermarket” that’s possible for the case of DPM I that does not require bundling. Consider a share of A purchased for \$5. The share is composed of $\$5 \cdot 1_A$ and $\$ \mathcal{P}_1 \cdot 1_A$. Now suppose the current price has moved from \$5 to \$10 per share and the trader wants to cash out at a profit. The trader can sell 1/2 share at market price (1/2 share for \$5), receiving all of the initial \$5 investment back, and retaining 1/2 share of A . The 1/2 share is worth either some positive amount, or nothing, depending on the outcome and the final payoff. So the trader is left with shares worth a positive expected value and all of his or her initial investment. The trader has essentially cashed out and locked in his or her gains. Now suppose instead that the price moves downward, from \$5 to \$2 per share. The trader decides to limit his or her loss by selling the share for \$2. The buyer gets the 1 share plus $\$2 \cdot 1_A$ (the buyer’s price refunded). The trader (seller) gets the \$2 plus what remains of the original price refunded, or $\$3 \cdot 1_A$. The trader’s loss is now limited to \$3 at most instead of \$5. If A occurs, the trader breaks even; if B occurs, the trader loses \$3.

Also note that—in either DPM formulation—traders can

always “hedge sell” by buying the opposite outcome without the need for any type of aftermarket.

8. CONCLUSIONS

I have presented a new market mechanism for wagering on, or hedging against, a future uncertain event, called a dynamic pari-mutuel market (DPM). The mechanism combines the infinite liquidity and risk-free nature of a pari-mutuel market with the dynamic nature of a CDA, making it suitable for continuous information aggregation. To my knowledge, all existing mechanisms—including the standard pari-mutuel market, the CDA, the CDAMMM, the bookie mechanism, and the MSR—exhibit at most two of the three properties. An MSR is the closest to a DPM in terms of these properties, if not in terms of mechanics. Given some natural constraints on price dynamics, I have derived in closed form the implied price functions, which encode how prices change continuously as shares are purchased. The interface for traders looks much like the familiar CDA, with the system acting as an automated market maker willing to accept an infinite number of buy orders at some price. I have explored two main variations of a DPM: one where only losing money is redistributed, and one where all money is redistributed. Each has its own pros and cons, and each supports several reasonable price functions. I have described the workings of an aftermarket, so that traders can cash out of the market early, like in a CDA, to lock in their gains or limit their losses, an operation that is not possible in a standard pari-mutuel setting.

9. FUTURE WORK

This paper reports the results of an initial investigation of the concept of a dynamic pari-mutuel market. Many avenues for future work present themselves, including the following:

- **Random walk conjecture.** The most important question mark in my mind is whether the random walk assumption (3) can be proven under reasonable market efficiency conditions and, if not, how severely it effects the practicality of the system.
- **Incentive analysis.** Formally, what are the incentives for traders to act on new information and when? How does the level of initial subsidy effect trader incentives?
- **Laboratory experiments and field tests.** This paper concentrated on the mathematics and algorithmics of the mechanism. However, the true test of the mechanism’s ability to serve as an instrument for hedging, wagering, or information aggregation is to test it with real traders in a realistic environment. In reality, how do people behave when faced with a DPM mechanism?
- **DPM call market.** I have derived the price functions to react to wagers on one outcome at a time. The mechanism could be generalized to accept orders on both sides, then update the prices wholistically, rather than by assuming a particular sequence on the wagers.
- **Real-valued variables.** I believe the mechanisms in this paper can easily be generalized to multiple discrete

outcomes, and multiple real-valued outcomes that always sum to some constant value (e.g., multiple percentage values that must sum to 100). However, the generalization to real-valued variables with arbitrary range is less clear, and open for future development.

- **Compound/combinatorial betting.** I believe that DPM may be well suited for compound [8, 11] or combinatorial [2] betting, for many of the same reasons that market scoring rules [11] are well suited for the task. DPM may also have some computational advantages over MSR, though this remains to be seen.

Acknowledgments

I thank Dan Fain, Gary Flake, Lance Fortnow, and Robin Hanson.

10. REFERENCES

- [1] Mukhtar M. Ali. Probability and utility estimates for racetrack bettors. *Journal of Political Economy*, 85(4):803–816, 1977.
- [2] Peter Bossaerts, Leslie Fine, and John Ledyard. Inducing liquidity in thin financial markets through combined-value trading mechanisms. *European Economic Review*, 46:1671–1695, 2002.
- [3] Kay-Yut Chen, Leslie R. Fine, and Bernardo A. Huberman. Forecasting uncertain events with small groups. In *Third ACM Conference on Electronic Commerce (EC’01)*, pages 58–64, 2001.
- [4] Sandip Debnath, David M. Pennock, C. Lee Giles, and Steve Lawrence. Information incorporation in online in-game sports betting markets. In *Fourth ACM Conference on Electronic Commerce (EC’03)*, 2003.
- [5] Robert Forsythe and Russell Lundholm. Information aggregation in an experimental market. *Econometrica*, 58(2):309–347, 1990.
- [6] Robert Forsythe, Forrest Nelson, George R. Neumann, and Jack Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.
- [7] Robert Forsythe, Thomas A. Rietz, and Thomas W. Ross. Wishes, expectations, and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39:83–110, 1999.
- [8] Lance Fortnow, Joe Kilian, David M. Pennock, and Michael P. Wellman. Betting boolean-style: A framework for trading in securities based on logical formulas. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce*, pages 144–155, 2003.
- [9] John M. Gandar, William H. Dare, Craig R. Brown, and Richard A. Zuber. Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance*, LIII(1):385–401, 1998.
- [10] Robin Hanson. Decision markets. *IEEE Intelligent Systems*, 14(3):16–19, 1999.
- [11] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1), 2002.
- [12] Robin D. Hanson. Could gambling save science? Encouraging an honest consensus. *Social Epistemology*, 9(1):3–33, 1995.
- [13] Jens Carsten Jackwerth and Mark Rubinstein. Recovering probability distributions from options prices. *Journal of Finance*, 51(5):1611–1631, 1996.
- [14] Joseph B. Kadane and Robert L. Winkler. Separating probability elicitation from utilities. *Journal of the American Statistical Association*, 83(402):357–363, 1988.
- [15] David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn Årup Nielsen. The real power of artificial markets. *Science*, 291:987–988, February 9 2001.
- [16] David M. Pennock, Steve Lawrence, Finn Årup Nielsen, and C. Lee Giles. Extracting collective probabilistic forecasts from web games. In *Seventh International Conference on Knowledge Discovery and Data Mining*, pages 174–183, 2001.
- [17] C. R. Plott, J. Wit, and W. C. Yang. Parimutuel betting markets as information aggregation devices: Experimental results. Technical Report Social Science Working Paper 986, California Institute of Technology, April 1997.
- [18] Charles R. Plott. Markets as information gathering tools. *Southern Economic Journal*, 67(1):1–15, 2000.
- [19] Charles R. Plott and Shyam Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.
- [20] Charles Polk, Robin Hanson, John Ledyard, and Takashi Ishikida. Policy analysis market: An electronic commerce application of a combinatorial information market. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce*, pages 272–273, 2003.
- [21] R. Roll. Orange juice and weather. *American Economic Review*, 74(5):861–880, 1984.
- [22] Richard N. Rosett. Gambling and rationality. *Journal of Political Economy*, 73(6):595–607, 1965.
- [23] Carsten Schmidt and Axel Werwatz. How accurate do markets predict the outcome of an event? The Euro 2000 soccer championships experiment. Technical Report 09-2002, Max Planck Institute for Research into Economic Systems, 2002.
- [24] Wayne W. Snyder. Horse racing: Testing the efficient markets model. *Journal of Finance*, 33(4):1109–1118, 1978.
- [25] Richard H. Thaler and William T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- [26] Martin Weitzman. Utility analysis and group behavior: An empirical study. *Journal of Political Economy*, 73(1):18–26, 1965.
- [27] Robert L. Winkler and Allan H. Murphy. Good probability assessors. *J. Applied Meteorology*, 7:751–758, 1968.