

An Axiomatic Characterization of Continuous-Outcome Market Makers^{*}

Xi Alice Gao and Yiling Chen

School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
{xagao,yiling}@eecs.harvard.edu

Abstract. Most existing market maker mechanisms for prediction markets are designed for events with a finite number of outcomes. All known attempts on designing market makers for forecasting continuous-outcome events resulted in mechanisms with undesirable properties. In this paper, we take an axiomatic approach to study whether it is possible for continuous-outcome market makers to satisfy certain desirable properties simultaneously. We define a general class of continuous-outcome market makers, which allows traders to express their information on any continuous subspace of their choice. We characterize desirable properties of these market makers using formal axioms. Our main result is an impossibility theorem showing that if a market maker offers binary-payoff contracts, either the market maker has unbounded worst case loss or the contract prices will stop being responsive, making future trades no longer profitable. In addition, we analyze a mechanism that does not belong to our framework. This mechanism has a worst case loss linear in the number of submitted orders, but encourages some undesirable strategic behavior.

Keywords: Prediction markets, continuous-outcome events, combinatorial prediction markets, expressive betting.

1 Introduction

A ubiquitous need in organizations and societies is to obtain and aggregate dispersed information of uncertain events so that informed decisions can be made. Prediction markets have been designed for this goal of information aggregation and have been shown to provide remarkably accurate forecasts in practice [1,2,3,4,5,6].

A prediction market is a betting intermediary that offers contracts whose payoffs are tied to outcomes of future events. Participants reveal their information about the event through buying and selling contracts. To facilitate information aggregation, many automated market maker mechanisms [7,8,9,10,11,12] have been designed to ensure that a participant can always conduct trades with the market maker and reveal his information whenever he finds it profitable.

Many events of interests, from carbon dioxide emission level to hurricane landing location, are naturally perceived as continuous random variables with

^{*} This work is supported by NSF under Grant No. CCF-0953516. Gao is partially supported by a NSERC PGS-D Scholarship.

continuous outcome spaces. However, most existing market makers, including the popular logarithmic market scoring rule (LMSR) [7,8] and the dynamic parimutuel markets (DPM) [9,13], are designed for finite-outcome random variables, and cannot handle continuous outcome spaces directly. For forecasting continuous-outcome events, these mechanisms rely on discretizing the continuous outcome space into a finite number of subsets and treat the event as a finite-outcome random variable. This approach poses the significant challenge of determining the level of discretization to be used in advance. Choosing too coarse-grained discretization could hurt information aggregation, since market participants may not be able to easily express their information with the pre-specified subsets. If the chosen discretization is too fine-grained, certain market makers like LMSR may suffer from a large worst case loss. In general, committing to an inappropriate discretization in advance may create unnecessary psychological burden for traders. In practice, Crowdcast¹, Yoopick [14], and Gates Hillman [15] prediction markets allow traders to wager on intervals through their user-friendly interfaces, although the underlying mechanisms still use some sort of discretization. Therefore, for predicting events that are naturally perceived as continuous, it is desirable to design market mechanisms that can handle the continuous-outcome spaces directly and provide sufficient expressiveness for participants to easily reveal their information on the continuous-outcome spaces.

Gao, Chen and Pennock [16] proposed the continuous-outcome LMSR and DPM. Although these continuous-outcome mechanisms offer considerable flexibility for participants to reveal their information, they suffer from some undesirable properties. In particular, the continuous-outcome LMSR can potentially lose an infinite amount of money, whereas the finite-outcome LMSR is guaranteed to have bounded worst case loss. The continuous DPM suffers from a different problem – even if a trader bets on a subspace that contains the realized outcome, he can potentially incur a loss. The intellectual quest that motivates this paper is to understand which set of desirable properties are possible or impossible to satisfy simultaneously for continuous-outcome market makers.

In this paper, we take an axiomatic approach to analyzing market makers for continuous-outcome events. We first define cost functional based market makers for continuous-outcome events. Then, we characterize desirable properties of these market makers as formal axioms. Our main contribution is an impossibility result showing that no market maker of this class can satisfy these axioms simultaneously. Specifically, for a market maker offering binary-payoff contracts, it either has unbounded worst-case loss or the contract prices will become unresponsive to trades being conducted. We also analyze a mechanism which does not fit into our axiomatic framework. This mechanism has a worst case loss linear in the number of orders, but encourages some undesirable strategic behavior.

Related Work. There have been a significant amount of efforts on designing and analyzing market maker mechanisms for finite-outcome events, including

¹ <http://www.crowdcast.com>

market scoring rules [7,8], dynamic parimutuel markets [9], cost function based market makers [10,17], and sequential convex parimutuel mechanisms [12]. The focus has been on analyzing the various properties of the market makers and establishing connections among them. In the context of designing combinatorial prediction markets, research has been focusing on the computational tractability of pricing expressive bets in the finite-outcome LMSR [18,19]. The work of Gao, Chen, and Pennock [16] is the closest to this paper. It generalized LMSR and DPM to handle continuous outcome spaces and analyzed the properties of the resulting mechanisms.

2 Background

In this section, we first describe a class of cost function based market makers for finite-outcome events. We then introduce and discuss the properties of the continuous-outcome LMSR market maker.

For finite-outcome events, Chen and Pennock [10] introduced a general class of automated market maker mechanisms, called the cost function based market makers. It has been shown that this class of market makers is equivalent to most of the known finite-outcome market makers under mild conditions [10,12,17]².

A cost function based market maker offers N contracts, each corresponding to one of N mutually exclusive and exhaustive outcomes of an event. Each contract pays off \$1 if and only if the corresponding outcome occurs. The market maker uses a differentiable *cost function* $C(\mathbf{q}) : \mathbb{R}^N \rightarrow \mathbb{R}$ to capture the total amount of money wagered in the market, where the vector \mathbf{q} represents the number of shares purchased by all traders. If a trader changes the quantity vector from \mathbf{q} to \mathbf{q}' , he pays $C(\mathbf{q}') - C(\mathbf{q})$ to the market maker and acquires $\mathbf{q}' - \mathbf{q}$ shares of contracts. The instantaneous price of the i -th contract, defined as $p_i(\mathbf{q}) = \partial C(\mathbf{q}) / \partial q_i$, represents the price per share of an infinitesimal number of shares.

Chen and Vaughan [17] formalized that a cost function is *valid* if the instantaneous prices $p_i(\mathbf{q})$ are non-negative and form a probability distribution over the outcome space. They proved that the sufficient and necessary conditions for a cost function C to be valid are: differentiability (to ensure that prices are well-defined), increasing monotonicity (to ensure that prices are non-negative), and a translation invariance condition $C(\mathbf{q} + k\mathbf{1}) = C(\mathbf{q}) + k$, $\forall \mathbf{q}, k$ (to ensure that prices sum to 1 and there is no arbitrage).

It has been shown that many valid market makers based on convex cost functions have bounded worst case loss [12,17]³, where the loss of the market maker is seen as a subsidy to promote information aggregation. For instance, the popular LMSR mechanism has bounded worst-case loss given by $b \log N$.

For continuous-outcome events, Gao, Chen and Pennock [16] generalized the finite-outcome LMSR for the *interval betting* setting. Even though the resulting continuous-outcome LMSR can handle interval bets for continuous-outcome

² DPM is an exception to this.

³ This is because a valid convex cost function based market maker is equivalent to a strictly proper market scoring rule under mild conditions. Any market scoring rule with a regular proper scoring rule has bounded worst case loss.

events, it suffers from unbounded loss – the market maker could potentially lose an infinite amount of money to the traders.

3 An Axiomatic Framework

In this section, we will define a general class of automated market maker mechanisms for continuous-outcome events, the *cost functional based market makers*. These market makers generalize the cost function based market makers for finite-outcome events to handle continuous-outcome spaces. We then propose three axioms to characterize desirable properties for these market makers.

3.1 Cost Functional Based Market Makers for Continuous-Outcome Events

Consider a continuous random variable X with domain $(L, U) = \{x : x \in \mathbb{R}, L \leq x \leq U, L \in \mathbb{R} \cup \{-\infty\}, U \in \mathbb{R} \cup \{+\infty\}\}$. Let $x \in (L, U)$ represent a particular outcome and let x^* denote the realized outcome in hindsight. We define a class of cost functional based market makers for predicting the realized value of X .

Cost functional based market makers are operated based on trading shares of contracts. First, we define the *quantity function* $q(x) \in \mathcal{L}^1(L, U)$, representing the number of shares purchased for outcome $x \in (L, U)$, which is analogous to the quantity vector \mathbf{q} in the finite-outcome case⁴. The value $q(x)$ can be thought as the total number of shares purchased for contracts that will pay off when x is the realized outcome.

A cost functional based market maker uses a differentiable *cost functional*, $C[q(x)] : \mathcal{L}^1(L, U) \rightarrow \mathbb{R}$, to capture the total amount of money wagered in the market as a functional of the current quantity function $q(x)$. If a trader changes the quantity function from $q(x)$ to $q'(x)$, he obtains $q'(x) - q(x)$ shares for each outcome x and must pay $C[q'(x)] - C[q(x)]$ to the market maker. We use $p[q(x), q'(x)]$ to denote the cost of such a transaction, i.e. $p[q(x), q'(x)] = C[q'(x)] - C[q(x)]$. The market maker starts the market with some initial quantity function $q^0(x)$ such that the value of $C[q^0(x)]$ is finite.

For any $q(x)$, the *price density functional* $p[q(x)]$ is defined as the functional derivative of the cost functional with respect to $q(x)$, that is, $p[q(x)] = \delta C[q(x)] / \delta q(x)$. The functional $p[q(x)]$ maps the quantity function to a probability density function over (L, U) . It is analogous to $p_i(\mathbf{q})$ in the finite-outcome setting. According to the calculus of functionals, we can express the cost of a transaction in terms of an integral of the price density functional, that is

$$\begin{aligned} p[q(x), q'(x)] &= C[q'(x)] - C[q(x)] \\ &= \int_0^1 \int_L^U p[q(x) + k(q'(x) - q(x))](q'(x) - q(x)) dx dk \end{aligned} \tag{1}$$

⁴ $\mathcal{L}^1(L, U)$ denotes the space of Lebesgue integrable functions on (L, U) with norm $\|q(x)\| = \int_L^U |q(x)| dx$.

If a trader changes the quantity function from $q(x)$ to $q'(x)$, then the future payoff of this transaction $o[q(x), q'(x), x^*]$ is a nonzero real number if $q'(x^*) \neq q(x^*)$ (i.e. the trader is buying or selling winning contracts), and \$0 otherwise where x^* is the realized outcome. Negative payoff encodes loss from selling the winning contracts. Other than this, we put no restriction on the value of $o[q(x), q'(x), x^*]$ and leave the definition of this value to specific mechanisms.

In our framework, a cost functional is *valid* if and only if the corresponding market maker satisfies two simple conditions:

1. For every $x \in (L, U)$, and every $q(x) \in \mathcal{L}^1(L, U)$, $p[q(x)] \geq 0$.
2. For every $q(x) \in \mathcal{L}^1(L, U)$, $\int_L^U p[q(x)]dx = 1$.

These are the minimum requirements for the price density functional to represent a valid probability distribution over the outcome space. The following theorem gives the sufficient and necessary conditions for the cost functional to be valid.

Theorem 1. *A cost functional C is valid if and only if it satisfies the following properties:*

1. *Differentiability:* The functional derivative $\delta C[q(x)]/\delta q(x)$ exists for all $q(x) \in \mathcal{L}^1(L, U)$ and all $x \in (L, U)$.
2. *Increasing Monotonicity:* For any $q(x), q'(x) \in \mathcal{L}^1(L, U)$, if $q'(x) \geq q(x), \forall x \in (L, U)$, then $C[q'(x)] \geq C[q(x)]$.
3. *Positive Translation Invariance:* For any $q(x) \in \mathcal{L}^1(L, U)$ and any constant k , $C[q(x) + k] = C[q(x)] + k$.

The above concepts define a general class of market maker mechanisms for forecasting continuous-outcome events⁵. These market makers can potentially support many different betting languages. In this paper, we focus on the simple and intuitive *interval betting* language [16]. For interval betting, traders are restricted to purchasing a constant s shares of a contract on an interval $(a, b) \subseteq (L, U)$ of their choice, where $a < b$. Such a transaction increases $q(x)$ by s for every $x \in (a, b)$. We denote the quantity function after the transaction by $q'(x) = \{q(x) + s\}_{(a,b)}$ where $q'(x)$ is defined by $q'(x) = q(x) + s, \forall x \in (a, b)$ and $q'(x) = q(x), \forall x \in (L, U) \setminus (a, b)$. For such a transaction, we define the *instantaneous contract price* $p_{(a,b)}[q(x)]$ to be the integral of the price density functional over (a, b) , that is, $p_{(a,b)}[q(x)] = \int_a^b p[q(x)]dx$. This is intuitively the price per share for buying an infinitesimal share of (a, b) . Note that we still do not put explicit restrictions on the transaction payoff for an interval contract (a, b) , i.e. $o[q(x), q'(x), x^*]$ where $q'(x) = \{q(x) + s\}_{(a,b)}$.

In the rest of the paper, we only consider cost functional based market makers for interval betting. Below we define an *Interval Cost Continuity* condition for interval betting market makers.

⁵ We note that the continuous-outcome DPM is not a valid market maker in our framework because its price density function does not correspond to a probability distribution.

Definition 2 (Interval Cost Continuity). A cost functional $C[q(x)]$ satisfies the Interval Cost Continuity condition if for any $x^* \in (L, U)$, $q(x) \in \mathcal{L}^1(L, U)$, $s \in \mathbb{R}$, and $q'(x) = \{q(x) + s\}_{(x^* - \delta, x^* + \delta)}$, $C[q'(x)]$ is right continuous at $\delta = 0$ and continuous for all $\delta > 0$.

The *Interval Cost Continuity* property specifies that, for each interval bet, the cost functional value for the final quantity function must be continuous for any change δ in the size of the betting interval. As δ approaches 0 (i.e. the size of the interval approaches 0), $C[q'(x)]$ approaches to $C[q(x)]$.

3.2 Desirable Properties

We propose three formal axioms to characterize some desirable properties of the cost functional based market makers for continuous-outcome events. We only consider interval bets.

Axiom 1 (Responsive Price). If $q(x)$ and $q'(x)$ satisfy three conditions: (1) $q'(x) = q(x), \forall x \in (L, U) \setminus (a, b)$, (2) $q'(x) \geq q(x), \forall x \in (a, b)$, and (3) $\exists(c, d) \subseteq (a, b)$ s.t. $q'(x) > q(x), \forall x \in (c, d)$, then

$$p_{(a,b)}[q'(x)] > p_{(a,b)}[q(x)]$$

for any contract $(a, b) \subset (L, U)$.

The *Responsive Price* axiom specifies that the instantaneous contract price is strictly monotonically increasing as the quantity over one of its subintervals strictly increases. This axiom is desirable since it guarantees that the change in the instantaneous contract prices will always *respond* to trades conducted and traders are always able to conduct trades irrespective of the current prices.

Axiom 2 (Domain Consistency). The payoff and cost of purchasing shares of (L, U) are always equal, that is, for all $q(x)$, $q'(x) = \{q(x) + s\}_{(L,U)}$, and $x^* \in (L, U)$, we have $o[q(x), q'(x), x^*] = p[q(x), q'(x)]$.

Intuitively, any bet on the entire domain (L, U) should earn zero profit as the bet is not revealing any useful information about X . This axiom is required for a cost functional based market maker to be arbitrage free. For instance, the continuous-outcome LMSR satisfies this axiom. Moreover, we will show in the next section that this axiom is a sufficient and necessary condition for the contracts offered to be exclusively binary-payoff contracts.

Axiom 3 (Bounded Loss). There exists $B \in \mathbb{R}$, such that, for any sequence of n transactions where the quantity functions satisfy $q_i(x) = \{q_{i-1}(x) + s_i\}_{(a_i, b_i)}$ and $(a_i, b_i) \subseteq (L, U)$, we have

$$\max_{x^* \in (L,U)} \left(\sum_{i=1}^n (o[q_i(x), q_{i+1}(x), x^*] - p[q_i(x), q_{i+1}(x)]) \right) \leq B.$$

This axiom gives a sufficient and necessary condition for the market maker to have bounded worst case loss. The market maker’s loss is the difference between the total money he pays out and the total money he collects. The worst outcome for the market maker is when this difference is maximized.

4 Impossibility Result

In this section, we present our main impossibility theorem. We first prove conditions for a valid market maker mechanism to offer exclusively binary-payoff contracts. For these market makers, we prove in Theorem 5 that the *Responsive Price* and *Bounded Loss* axioms cannot be satisfied simultaneously.

Lemma 3 (Binary Contract Lemma). *A valid market maker mechanism satisfies the Domain Consistency axiom if and only if it offers binary-payoff interval contracts, that is, the future payoff of any contract is fixed to be \$1 per share if the realized outcome x^* falls within the interval and \$0 otherwise.*

Lemma 3 shows that if a valid market maker satisfies the *Domain Consistency* axiom, the payoff of the contract has to be binary regardless of the interval chosen. We also note that with binary-payoff contracts, the *Responsive Price* axiom implies that the price of any contract never reaches 0 or 1. Next, we present Lemma 4 to facilitate the proof of our main impossibility result.

Lemma 4 (Responsive Price Lemma). *For a valid market maker satisfying the Interval Cost Continuity condition, if it satisfies the Responsive Price axiom, then for any winning contract (a, b) , any number of shares $s \in \mathbb{Z}^+$, any quantity function $q(x)$, and any $\epsilon > 0$, there exists a winning contract $(a', b') \subset (a, b)$, such that*

$$C[q''(x)] - C[q(x)] \leq \epsilon(C[q'(x)] - C[q(x)])$$

where $q'(x) = \{q(x) + s\}_{(a,b)}$ and $q''(x) = \{q(x) + s\}_{(a',b')}$.

Based on Lemma 4, if an interval (a, b) is a winning contract, there exists a subinterval of (a, b) which is also winning such that the cost of buying a constant number of shares over the subinterval is an arbitrarily small fraction of the cost of buying the same number of shares over (a, b) .

Theorem 5 (Impossibility Result). *For a valid market maker satisfying the Interval Cost Continuity condition, if it allows traders to bet on intervals of any nonzero size and satisfies the Domain Consistency axiom, then it cannot satisfy the Responsive Price and Bounded Loss axioms simultaneously.*

Proof Sketch. By Lemma 3, the contracts offered must pay off \$1 per share if they are winning, and \$0 otherwise. Consider a trader who knows x^* and has a fixed budget of \$ m . Using the following procedure, this trader could potentially get an arbitrarily large profit.

To get \$ s payoff, the trader can start by calculating the cost of buying s shares of an arbitrary winning contract (a, b) , denoted by T . If $T \geq m$, then by Lemma 4, the trader can choose $\epsilon = \frac{m}{T}$ and find a winning contract $(a', b') \subset (a, b)$ such that the cost of buying s shares over (a', b') is no more than $\frac{m}{T} \cdot T = m$ dollars and the corresponding profit is at least $s - m$ dollars. Because s is arbitrary, the trader's profit, hence the market maker's loss, is not bounded. \square

Even though Theorem 5 allows traders to bet on intervals of any nonzero size, we now show that even if we restrict the size of the smallest betting interval to be

at least $z > 0$, with certain assumptions, the trader could still bet on arbitrarily small intervals of their choice.

Corollary 6. *For a valid market maker satisfying the Interval Cost Continuity condition and restricting the size of the smallest betting interval to be $z \in \mathbb{R}$ where $0 < z < (U - L)/2^6$, if it satisfies the Domain Consistency axiom, then it cannot simultaneously satisfy the Responsive Price and Bounded Loss axioms.*

The key insight for proving the above corollary is that a trader can perform a sequence of transactions which is equivalent to purchasing shares of an arbitrarily small interval even with the restriction on the size of the smallest betting interval. It is worth noting that these transactions can be potentially completed in any order, and multiple traders can collude to complete them. Thus, it would be very challenging in general to detect such trading patterns in practice.

The unbounded worst case loss of the continuous-outcome LMSR is a special case of our impossibility result. However, compared with finite-outcome market makers, this impossibility result is rather surprising since the finite-outcome LMSR essentially satisfies the finite-outcome versions of all three axioms.

We could possibly relax the *Responsive Price* axiom to derive mechanisms with bounded worst case loss, although the resulting market maker may be trivial and less interesting. For example, a market maker can quickly increase the price of contracts to 1 once the quantity for the contracts increases beyond a certain value. Beyond this point, purchasing more shares will not earn the trader any more profit and bounded worst case loss can be achieved.

5 Discussion and Conclusion

While the class of market makers we considered is quite general, there exist other continuous-outcome mechanisms outside of this class that can achieve bounded worst case loss and responsive price simultaneously. In particular, we can operate the finite-outcome LMSR over the continuous-outcome space by discretizing the outcome space on the fly given the submitted orders. By violating our definition of *instantaneous contract price* and the *Interval Cost Continuity* condition, this mechanism achieves the worst case loss linear in the number of orders submitted, but also encourages certain undesirable strategic behaviors.

To operate the finite-outcome LMSR over a continuous-outcome space, we split the existing intervals in the state space for every submitted order on (c, d) whenever c or d falls within the existing intervals. Then the order is traded via LMSR with the state space after splitting. For every splitting of this kind, the outstanding quantities and instantaneous prices need to satisfy the following consistency constraints for the mechanism to remain arbitrage free.

- The sum of the instantaneous prices of the subintervals must be equal to the instantaneous price of the original interval.

⁶ This assumption is reasonable since the size of the smallest betting interval should be much smaller than the size of the domain of the random variable.

- The number of shares of each subinterval held by all traders must be equal to the number of shares of the original interval held by all traders.

The first consistency constraint allows considerable freedom in splitting the probability estimates among the subintervals. If the probability estimates are split equally among the subintervals, then the resulting mechanism violates our definition of *instantaneous contract price* and the *Interval Cost Continuity* condition. However, this market maker has worst case loss given by $M \log 3$, where M is the number of orders submitted. However, this mechanism does not provide the incentive for traders to reveal their information truthfully. Given several subintervals with equal prices, a trader could maximize his probability of winning by betting on the largest interval regardless of his subject probability estimates for these intervals.

If the market maker splits the probability estimates in proportion to the lengths of the subintervals, it satisfies all the axioms proposed and the worst case loss becomes unbounded according to Theorem 5. Intuitively, the unbounded loss is due to the market maker assigning arbitrarily small initial probability to the smallest interval containing the realized outcome. If the traders drive the price of this interval to be \$1, then the market maker is destined to lose an infinite amount of money to the traders.

In conclusion, we take an axiomatic approach to study automated market maker mechanisms for forecasting continuous-outcome events. In our axiomatic framework, we consider a general class of cost functional based market makers and define formal axioms to characterize desirable properties of these mechanisms. We then prove that it is impossible for a valid cost functional based market maker mechanism to satisfy a certain set of properties simultaneously. Our results suggest that future efforts on designing continuous-outcome market makers should focus on finding the right tradeoffs among the desirable properties. In particular, it may be fruitful to investigate whether a market maker can be designed such that the *Interval Cost Continuity* condition or the *Responsive Price* axiom is relaxed to a reasonable degree so that other desirable properties can be achieved for practical applications.

References

1. Berg, J.E., Forsythe, R., Nelson, F.D., Rietz, T.A.: Results from a dozen years of election futures markets research. In: Plott, C.A., Smith, V. (eds.) *Handbook of Experimental Economic Results* (2001)
2. Wolfers, J., Zitzewitz, E.: Prediction markets. *Journal of Economic Perspective* 18(2), 107–126 (2004)
3. Forsythe, R., Nelson, F., Neumann, G.R., Wright, J.: Anatomy of an experimental political stock market. *American Economic Review* 82(5), 1142–1161 (1992)
4. Forsythe, R., Rietz, T.A., Ross, T.W.: Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior & Organization* 39(1), 83–110 (1999)
5. Debnath, S., Pennock, D.M., Giles, C.L., Lawrence, S.: Information incorporation in online in-game sports betting markets. In: *EC 2003: Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 258–259. ACM, New York (2003)

6. Chen, K.Y., Plott, C.R.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working Paper No.1131, California Institute of Technology (2002)
7. Hanson, R.D.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
8. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* 1(1), 1–15 (2007)
9. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *EC 2004: Proceedings of the 5th ACM Conference on Electronic Commerce*, pp. 170–179. ACM, New York (2004)
10. Chen, Y., Pennock, D.M.: A utility framework for bounded-loss market makers. In: *UAI 2007: Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pp. 49–56 (2007)
11. Peters, M., So, A.M.C., Ye, Y.: Pari-mutuel markets: Mechanisms and performance. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 82–95. Springer, Heidelberg (2007)
12. Agrawal, S., Delage, E., Peters, M., Wang, Z., Ye, Y.: A unified framework for dynamic pari-mutuel information market design. In: *EC 2009: Proceedings of the Tenth ACM Conference on Electronic Commerce*, pp. 255–264. ACM, New York (2009)
13. Chen, Y., Pennock, D.M., Kasturi, T.: An empirical study of dynamic pari-mutuel markets: Evidence from the Tech Buzz Game. In: *The Workshop on Web Mining and Web Usage Analysis, WebKDD*. LNCS, Springer, Heidelberg (2008)
14. Goel, S., Pennock, D., Reeves, D.M., Yu, C.: Yoopick: A combinatorial sports prediction market. In: *AAAI*, pp. 1880–1881 (2008)
15. Othman, A., Sandholm, T.: Automated market-making in the large: the gates hillman prediction market. In: *EC 2010: Proceedings of the 11th ACM Conference on Electronic Commerce*, pp. 367–376. ACM, New York (2010)
16. Gao, X., Chen, Y., Pennock, D.M.: Betting on the real line. In: Leonardi, S. (ed.) *WINE 2009*. LNCS, vol. 5929, pp. 553–560. Springer, Heidelberg (2009)
17. Chen, Y., Vaughan, J.W.: A new understanding of prediction markets via no-regret learning. In: *EC 2010: Proceedings of the 11th ACM Conference on Electronic Commerce*, pp. 189–198. ACM, New York (2010)
18. Chen, Y., Fortnow, L., Lambert, N., Pennock, D.M., Wortman, J.: Complexity of combinatorial market makers. In: *EC 2008: Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 190–199. ACM, New York (2008)
19. Chen, Y., Goel, S., Pennock, D.M.: Pricing combinatorial markets for tournaments. In: *STOC 2008: Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pp. 305–314. ACM, New York (2008)