

Liquidity-Sensitive Automated Market Makers via Homogeneous Risk Measures

Abraham Othman and Tuomas Sandholm

Computer Science Department, Carnegie Mellon University
{aothman,sandholm}@cs.cmu.edu

Abstract. Automated market makers are algorithmic agents that provide liquidity in electronic markets. A recent stream of research in automated market making is the design of *liquidity-sensitive* automated market makers, which are able to adjust their price response to the level of active interest in the market. In this paper, we introduce homogeneous risk measures, the general class of liquidity-sensitive automated market makers, and show that members of this class are (necessarily and sufficiently) the convex conjugates of compact convex sets in the non-negative orthant. We discuss the relation between features of this convex conjugate set and features of the corresponding automated market maker in detail, and prove that it is the curvature of the convex conjugate set that is responsible for implicitly regularizing the price response of the market maker. We use our insights into the dual space to develop a new family of liquidity-sensitive automated market makers with desirable properties.

1 Introduction

Automated market makers are algorithmic agents that provide liquidity in electronic markets. Markets with large event spaces or sparse interest from traders might fail because buyers and sellers have trouble finding one another. Automated market makers can prevent this failure by stepping in and providing a counterparty for prospective traders; instead of making bets with each other, traders place bets with the automated market maker. Automated market makers have been the object of theoretical study into market microstructure [Ostrovsky, 2009, Othman and Sandholm, 2010b] and successfully implemented in practice in large electronic markets [Goel et al., 2008, Othman and Sandholm, 2010a]. A broad introduction to the mechanics of automated market making can be found in Pennock and Sami [2007].

Othman et al. [2010] introduce a *liquidity-sensitive* automated market maker. This market maker is able to adapt its price response to increasing activity within the market; with this market maker bets will not move prices very much when there is lots of money already wagered with the market maker. This is in contrast to traditional market-making agents that provide identical price responses regardless of whether there are tens of dollars or tens of millions of dollars wagered with the market maker.

Unfortunately, this liquidity-sensitive market maker does not generalize easily. In Othman et al. [2010] it is referred to as a technique to “continuously channel profits into liquidity”, a view echoed by Abernethy et al. [2011]. While this view may be accurate, it is not prescriptive: it offers no insight about how to create other liquidity-sensitive market makers, or of the relation between liquidity-sensitive market makers and the other market makers of the literature.

In this paper, we solve the puzzle of how liquidity-sensitive market makers work, and their relation to other market makers from the literature. We are able to contextualize, generalize, and expand the idea of liquidity-sensitive market makers. In order to do this, we first situate liquidity-sensitive market makers within the same framework as their liquidity-insensitive counterparts. Using a set of desiderata taken from the prediction market and finance literature we introduce a new class of automated market makers, *homogeneous risk measures*, which we argue correctly embody the notion of liquidity sensitivity, and we prove that the market maker of Othman et al. [2010] is a member of this class.

Our principal result is a necessary and sufficient characterization of the complete set of homogeneous risk measures: they are the support functions of compact convex sets in the non-negative orthant. Most intriguingly, this dual view allows us to achieve a synthesis between homogeneous risk measures and the experts algorithm perspective of Chen and Vaughan [2010], another recent view of automated market making. In this perspective, homogeneous risk measures are *unregularized* follow-the-leader algorithms that (generally) put non-unit total weight on the set of experts. We show it is the shape of the convex conjugate set (particularly, that set’s curvature) that implicitly acts as a regularizer for the homogeneous risk measure. Furthermore, the bulge of the convex set away from the probability simplex defines notions like the maximum sum of prices. We use these insights to create a new family of liquidity-sensitive automated market makers, the *unit ball market makers*, that have desirable properties: defined costs for any possible bet, defined bounds on sums of prices, and tightly bounded loss.

2 Background

In this section we provide a brief introduction to automated market making, with emphasis on the recent results that guide the remainder of the work.

2.1 Cost functions and risk measures

We consider a general setting in which the future state of the world is exhaustively partitioned into n events, $\{\omega_1, \dots, \omega_n\}$, so that exactly one of the ω_i will occur. This model applies to a wide variety of settings, including financial models on stock prices and interest rates, sports betting, and traditional prediction markets.

In our notation, \mathbf{x} is a vector and x is a scalar, $\mathbf{1}$ is the n -dimensional vector of all ones, and $\nabla_i f$ represents the i -th element of the gradient of a function f . The non-negative orthant is given by $\mathbb{R}_+^n \equiv \{\mathbf{x} \mid \min_i x_i \geq 0\}$.

Let U be a convex subset of \mathbb{R}^n . Our work concerns functions $C : U \mapsto \mathbb{R}$ which map vector payouts over the events to scalar values. A *state* refers to a vector of payouts. Traders make bets with the market maker by changing the market maker’s state. To move the market maker from state \mathbf{x} to state \mathbf{x}' , traders pay $C(\mathbf{x}') - C(\mathbf{x})$. For instance, if the state is $x_1 = 5$ and $x_2 = 3$, then the market maker needs to pay out five dollars if ω_1 is realized and pay out 3 dollars if ω_2 is realized. If a new trader wants a bet that pays out one dollar if event ω_1 occurs, then they change the market maker’s state to be $\{6, 3\}$, and pay $C(\{6, 3\}) - C(\{5, 3\})$. There are two broad research streams that explore these functions. The prediction market literature, where they are called *cost functions*, and the finance literature, where they are called *risk measures*. We use the terms *cost function* and *risk measure* interchangeably.

The most popular cost function used in Internet prediction markets is Hanson’s logarithmic market scoring rule (LMSR), an automated market maker with particularly desirable properties, including bounded loss and a simple analytical form [Hanson, 2003, 2007]. The LMSR is defined as

$$C(\mathbf{x}) = b \log \left(\sum_i \exp(x_i/b) \right)$$

for fixed $b > 0$. b is called the *liquidity parameter*, because it controls the magnitude of the price response of the market maker to bets.¹ For instance, if the LMSR is used with $b = 10$ in our example above, $C(\{6, 3\}) - C(\{5, 3\}) \approx .56$, and so the market maker would quote a price of 56 cents to the agent for their bet. If $b = 1$, the same bet would cost 92 cents.

The *prices* p_i of a differentiable risk measure are given by the gradient of the cost function—the marginal cost on each event: $p_i = \frac{\exp(x_i/b)}{\sum_j \exp(x_j/b)}$. Observe that the prices in the LMSR sum to one. The notion of *sum of prices* is crucial to our work. The market maker’s profit cut (or *vigorish* in gambling contexts) can be thought of as the difference between the sum of prices and unity [Othman et al., 2010]. This profit cut serves to compensate the market maker for taking bets with traders, and typical values for the vigorish in real applications are small, ranging from one percent to 20 percent. Since the LMSR and many other cost functions of the literature [Chen and Pennock, 2007, Peters et al., 2007, Agrawal et al., 2009, Abernethy et al., 2011] do not have a profit cut, they can be expected to run at a loss in practice [Pennock and Sami, 2007].

2.2 Link to online learning

One of the most intriguing recent developments in automated market making is the link between cost functions and online learning algorithms, particularly between cost functions and online follow-the-regularized-leader algorithms. This link first appeared in a supporting role in Chen et al. [2008], and was significantly

¹ With $b = 1$, the LMSR is equivalent to the *entropic risk measure* of the finance literature [Föllmer and Schied, 2002].

expanded in later work by those authors [Chen and Vaughan, 2010, Abernethy et al., 2011]. Any loss-bounded convex risk measure (Section 3 will make this precise) is equivalent to a no-regret follow-the-regularized-leader online learning algorithm. These online learning algorithms are conventionally expressed not as cost functions (or, in the machine learning literature, *potential functions*), but rather in dual space [Shalev-Shwartz and Singer, 2007]. The dual-space formulation is a powerful way of interpreting and constructing automated market makers.

Let Π be the probability simplex. Chen and Vaughan [2010] show that we can write any convex risk measure in terms of a convex optimization over a *follow-the-leader* term and a convex *regularizer* term. This optimization is in fact a conjugacy operation restricted to the probability simplex: $C(\mathbf{x}) = \max_{\mathbf{y} \in \Pi} \mathbf{x} \cdot \mathbf{y} - f(\mathbf{y})$. Here, $\mathbf{x} \cdot \mathbf{y}$ is the follow-the-leader term, and f is a regularizer.

2.3 The OPRS cost function

The *Othman-Pennock-Reeves-Sandholm cost function (OPRS)* was originally introduced in Othman et al. [2010] as a liquidity-sensitive extension of the LMSR. The OPRS is defined as $C(\mathbf{x}) = b(\mathbf{x}) \log(\sum_i \exp(x_i/b(\mathbf{x})))$, where $b(\mathbf{x}) = \alpha \sum_i x_i$ for $\alpha > 0$. The OPRS can be contrasted with the LMSR, for which $b(\mathbf{x}) \equiv b$. Unlike the LMSR, the OPRS is only defined over the non-negative orthant (for continuity we can set $C(\mathbf{0}) = 0$). Also unlike the LMSR, the sum of prices in the OPRS is always greater than 1.

The OPRS has several desirable properties. These include a concise analytical closed form and *outcome-independent profit*, the ability to (for certain final quantity vectors) book a profit regardless of the realized outcome. Perhaps the most practical property of the OPRS is its scale-invariant liquidity sensitivity: its consistent price reaction over different scales of market activity. (This scale-invariance is a consequence of the OPRS cost function being positive homogeneous.) For large liquid markets, say with millions of dollars, a one-dollar bet will have a much smaller impact on prices than in a less-liquid market. This is not the case for the LMSR, where a one dollar bet moves prices the same amount in both heavily- and lightly-traded markets.

3 Desiderata, dual spaces, and an impossibility result

This section expands upon the dual-space approach to automated market making [Agrawal et al., 2009, Chen and Vaughan, 2010, Abernethy et al., 2011], particularly as a vehicle for contextualizing and generalizing the OPRS.

3.1 Desiderata and their combinations

In this section we introduce five desiderata for cost functions. Each of these properties has been acknowledged as desirable in the market making literature [Agrawal et al., 2009, Othman et al., 2010, Abernethy et al., 2011]. The market makers from the literature satisfy various subsets of these desiderata.

Monotonicity: For all \mathbf{x} and \mathbf{y} such that $x_i \leq y_i$, $C(\mathbf{x}) \leq C(\mathbf{y})$.

Monotonicity prevents simple arbitrages like a trader buying a zero-cost contract that never results in losses but sometimes results in gains.

Convexity: For all \mathbf{x} and \mathbf{y} and $\lambda \in [0, 1]$, $C(\lambda\mathbf{x} + (1 - \lambda)\mathbf{y}) \leq \lambda C(\mathbf{x}) + (1 - \lambda)C(\mathbf{y})$.

Convexity can be thought of as a condition that encourages diversification. The cost of the blend of two payout vectors is not greater than the sum of the cost of each individually. Consequently, the market maker is incentivized to diversify away its risk. The acknowledgment of diversification as desirable goes back to the very beginning of the mathematical finance literature [Markowitz, 1952].

Bounded loss: $\sup_{\mathbf{x}} [\max_i (x_i) - C(\mathbf{x})] < \infty$.

A market maker using a cost function with bounded loss can only lose a finite amount to interacting traders, regardless of the traders' actions and the realized outcome.

Translation invariance: For all \mathbf{x} and scalar α , $C(\mathbf{x} + \alpha\mathbf{1}) = C(\mathbf{x}) + \alpha$.

Translation invariance ensures that adding a dollar to the payout of every state of the world will cost a dollar.

Positive homogeneity: For all \mathbf{x} and scalar $\gamma > 0$, $C(\gamma\mathbf{x}) = \gamma C(\mathbf{x})$.

Positive homogeneity ensures a scale-invariant, currency-independent price response, as in the OPRS. From a risk measurement perspective, positive homogeneity ensures that doubling a risk doubles its cost.

A cost function that satisfies all of these desiderata except bounded loss is called a *coherent risk measure*. Coherent risk measures were first introduced in Artzner et al. [1999].

Definition 1 A coherent risk measure is a cost function that satisfies monotonicity, convexity, translation invariance, and positive homogeneity.

When we relax positive homogeneity from a coherent risk measure, we get a convex risk measure. Convex risk measures were first introduced in Carr et al. [2001] and feature prominently in the prediction market literature [Hanson, 2003, Ben-Tal and Teboulle, 2007, Hanson, 2007, Chen and Pennock, 2007, Peters et al., 2007, Agrawal et al., 2009, Abernethy et al., 2011].

When we instead relax translation invariance from a coherent risk measure, we get what we dub a homogeneous risk measure.

Definition 2 A homogeneous risk measure is a cost function satisfying monotonicity, convexity, and positive homogeneity.

To our knowledge, the only homogeneous risk measure of the literature that is not also a coherent risk measure is the OPRS.

Proposition 1 The OPRS is a homogeneous risk measure (for vectors in the non-negative orthant).

The desiderata are global properties that need to hold over the entire space the cost function is defined over. It is often difficult to verify that a given cost function satisfies these desiderata directly, and inversely, it is difficult to construct new cost functions that satisfy specific desiderata. Remarkably, each of these desiderata have simple representations in Legendre-Fenchel dual space.

3.2 Dual space equivalences

The rest of the paper relies on the well-developed theory of convex conjugacy.

Definition 3 *The Legendre-Fenchel dual (aka convex conjugate) of a convex cost function C is a convex function $f : \mathbb{Y} \mapsto \mathbb{R}$ over a convex set $\mathbb{Y} \subset \mathbb{R}^n$ such that $C(\mathbf{x}) = \max_{\mathbf{y} \in \mathbb{Y}} [\mathbf{x} \cdot \mathbf{y} - f(\mathbf{y})]$. We say that the cost function is “conjugate to” the pair \mathbb{Y} and f . Convex conjugates exist uniquely for convex cost functions defined over \mathbb{R}^n [Rockafellar, 1970, Boyd and Vandenberghe, 2004].*

We will refer to the convex optimization in dual space as the “optimization” or “optimization problem”, and the maximizing \mathbf{y} as the “maximizing argument”. One way of interpreting the dual is that it represents the “price space” of the market maker, as opposed to a cost function which is defined over a “quantity space” [Abernethy et al., 2011]. The only prices a market maker can assume are those $\mathbf{y} \in \mathbb{Y}$, while the function f serves as a measure of market sensitivity and a way to limit how quickly prices are adjusted in response to bets. As we have discussed, in the prediction market literature “prices” denote the partial derivatives of the cost function [Pennock and Sami, 2007, Othman et al., 2010]. When it is unique, the maximizing argument of the convex conjugate is the gradient of the cost function, and when it is not unique, then the maximizing arguments represent the subgradients of the cost function. A fuller discussion of the relation between convex conjugates and derivatives is available in convex analysis texts [Rockafellar, 1970, Boyd and Vandenberghe, 2004].

Another interpretation of the dual space is from online learning, specifically online regularized follow-the-leader algorithms [Chen and Vaughan, 2010]. We discussed the literature relating to this link in Section 2.2. Here, the set \mathbb{Y} represents the allowable weights we can assign to experts, and the function f is a regularizer that determines how quickly we adjust the weight between experts in response to returns which are the same as payouts in this interpretation. Generally speaking when the set \mathbb{Y} exceeds the probability simplex \mathcal{I} , then the weights placed on the experts will not be guaranteed to sum to unity.

With these interpretations in mind, we proceed to show the power of the dual space: we can represent homogeneous risk measures with a compact convex set in the non-negative orthant. The relations between convex and monotonic cost functions, convex and positive homogeneous cost functions, and their respective duals are a consequence of well-known results in the convex analysis literature [Rockafellar, 1966, 1970].

Proposition 2 *A risk measure is convex and monotonic if and only if the set \mathbb{Y} is exclusively within the non-negative orthant.*

Proposition 3 *A risk measure is convex and positive homogeneous if and only if its convex conjugate has compact \mathbb{Y} and has $f(\mathbf{y}) = 0$ for every $\mathbf{y} \in \mathbb{Y}$.*

In the literature this latter result relates *indicator sets* (here, the set \mathbb{Y}) to *support functions* (here, the cost function). Since $f(\mathbf{y}) = 0$ for all $\mathbf{y} \in \mathbb{Y}$, the cost function conjugacy is defined only by the set \mathbb{Y} . Consequently, we will abuse terminology slightly and refer to the cost functions as conjugate to the convex compact set alone. A necessary and sufficient condition on the set of homogeneous risk measures follows.

Corollary 1 *A cost function is a homogeneous risk measure if and only if it is conjugate to a compact convex set in the non-negative orthant.*

The following results can be derived from convex analysis and the work of Abernethy et al. [2011].

Proposition 4 *A risk measure is convex, monotonic, and translation invariant if and only if the set \mathbb{Y} lies exclusively on the probability simplex.*

Proposition 5 *A risk measure is convex and has bounded loss if and only if the set \mathbb{Y} includes the probability simplex.*

The only market maker that satisfies all five of our desiderata is max.

Proposition 6 *The only coherent risk measure with bounded loss is $C(\mathbf{x}) = \max_i x_i$.*

The max market maker corresponds to an order-matching, risk-averse cost function that either charges agents nothing for their transactions, or exactly as much as they could be expected to gain in the best case. For instance, a trader wishing to move the max market maker from state $\{5, 3\}$ to state $\{7, 3\}$ would be charged 2 dollars, exactly as much as they would win if the first event happened—which means taking the bet is a dominated action. On the other hand, a trader wishing to move the market maker from state $\{5, 3\}$ to state $\{5, 5\}$ pays nothing! These two small examples suggest that max is a poor risk measure in practice, and therefore Proposition 6 should be viewed as an impossibility result.

Combining all of our dual-space equivalences, we have that the conjugate of max is defined exclusively on the whole probability simplex, where it is identically 0.²

There are now two ways to smooth out the price response of max: One way is to use a regularizer, so that price estimates do not immediately jump to the axes (i.e., zero or one). This corresponds to a regularized online follow-the-leader algorithm, which is a convex risk measure [Chen and Vaughan, 2010]. We introduce a different approach, to expand the shape of the valid price, so that the shape of the space itself serves as an implicit regularizer over the price estimates. This will generally result in prices that are not probability distributions, and as we explore in the next section, this approach leads to homogeneous risk measures.

² In dual price space, the maximizing argument to the max cost function can always be represented as one of the axes. In the online learning view, max represents an unregularized follow-the-leader algorithm, putting all of its probability weight on the current best expert (i.e., the event with largest current payout).

4 Shaping the dual space

Recall that only the convex conjugate set \mathbb{Y} of a homogenous risk measure is responsible for determining the market maker's behavior, because the conjugate function f takes value zero everywhere in that set. In this section, we explore two features of the conjugate set that produce desirable properties: its *curvature* and its *divergence* from the probability simplex.

4.1 Curvature

We would like for our cost function to always be differentiable (outside of $\mathbf{0}$, where a derivative of a positive homogeneous function will not generally exist). The OPRS is differentiable in the non-negative orthant (again, excepting $\mathbf{0}$) while max is differentiable only when the maximum is unique. In this section, we show that only curved conjugate sets produce homogeneous risk measures that are differentiable.³

Definition 4 *A closed, convex set \mathbb{Y} is strictly convex if its boundary does not contain a non-degenerate line segment. Formally, let $\partial\mathbb{Y}$ denote the boundary of the set. Let $0 \leq \lambda \leq 1$ and $\mathbf{x}, \mathbf{x}' \in \partial\mathbb{Y}$. Then $\lambda\mathbf{x} + (1 - \lambda)\mathbf{x}' \in \partial\mathbb{Y}$ holds only for $\mathbf{x} = \mathbf{x}'$.*

Since strictly convex sets are never linear on their boundary they can be thought of as sets with curved boundaries.

Proposition 7 *A homogeneous risk measure is differentiable on $\mathbb{R}^n \setminus \mathbf{0}$ if and only if its conjugate set is strictly convex.*

4.2 Divergence from probability simplex

The amount of divergence from the probability simplex governs the market maker's divergence from translation-invariant prices (i.e., prices that sum to unity). Recall that max is the homogeneous risk measure that is defined only over the probability simplex.

Proposition 8 *Let \mathbb{Y} be the dual set of a differentiable homogeneous risk measure. Then the maximum sum of prices (the most a trader would ever need to spend for a unit guaranteed payout) is given by $\max_{\mathbf{y} \in \mathbb{Y}} \sum_i y_i$, and the minimum sum of prices (the most the market maker would ever pay for a unit guaranteed payout) is given by $\min_{\mathbf{y} \in \mathbb{Y}} \sum_i y_i$.*

³ It might be argued that what we are really interested in, particularly if we claim that curved sets act as a regularizer in the price response, is whether or not curved sets also imply *continuous* differentiability of the cost function. Continuous differentiability would mean that prices both exist and are continuous in the quantity vector. These conditions are in fact the same for convex functions defined over an open interval (such as $\mathbb{R}^n \setminus \mathbf{0}$), because for such functions differentiability implies continuous differentiability [Rockafellar, 1970].

Given any (efficiently representable) convex set corresponding to a differentiable homogeneous risk measure, the extreme price sums can be solved for in polynomial time, since it is a convex optimization over a convex set.

It was shown in Othman et al. [2010] that the OPRS achieved its maximum sum of prices for quantity vectors that are scalar multiples of $\mathbf{1}$. A corollary of the above result is that this property holds for every homogeneous risk measure. (Other vectors may also achieve the same sum of prices.)

Corollary 2 *In a homogeneous risk measure every vector that is a positive multiple of $\mathbf{1}$ achieves the maximum sum of prices.*

In addition to maximum prices, the shape of the convex set also determines the worst-case loss of the resulting market maker. The notion of worst-case loss is closely related to our desideratum of bounded loss—a market maker with unbounded worst-case loss does not have bounded loss, and a market maker with finite worst-case loss has bounded loss.

Definition 5 *The worst-case loss of a market maker is given by $\max_i x_i - C(\mathbf{x}) + C(\mathbf{x}^0)$ where $\mathbf{x}^0 \in \mathbb{R}_+^n$ is some initial quantity vector the market maker selects.*

In homogeneous risk measures, the amount of liquidity sensitivity is proportional to the market’s state. Since in practice there is some latent level of interest in trading on the event before the market’s initiation, it is desirable to seed the market initially to reflect a certain level of liquidity. It is desirable to have a tight bound on that worst-case loss, reflecting that in practice, market administrators are likely to have bounds on how much the market maker could lose in the worst case. Tight bounds on worst-case loss assure the administrator that that bound will be satisfied with maximum liquidity injected at the market’s initiation.

Proposition 9 *Let \mathbb{Y} be a convex set conjugate to a homogeneous risk measure that includes the unit axes but does not exceed the unit hypercube. Then the worst-case loss of the risk measure is tightly bounded by the initial cost of the market’s starting point.*

By bringing \mathbf{x}^0 as close as desired to $\mathbf{0}$, we have the following corollary, which is a generalization of a similar result for the OPRS.

Corollary 3 *Let \mathbb{Y} be a convex set conjugate to a homogeneous risk measure that includes the unit axes. Then the worst-case loss of the risk measure can be set arbitrarily small.*

A bound on prices also emerges from this result.

Corollary 4 *Let \mathbb{Y} be a convex set conjugate to a homogeneous risk measure that includes the unit axes but does not exceed the unit hypercube. Then the maximum price on any event is 1.*

5 A new family of liquidity-sensitive market makers

We proceed to use our theoretical results constructively, to create a family of homogeneous risk measures with desirable properties that the OPRS, the only prior homogeneous risk measure, lacks. These include tight bounds on minimum sum of prices and worst-case losses, and definition over all of \mathbb{R}^n . Our new family of market makers is parameterized (in much the same way as the OPRS) by the maximum sum of prices. The OPRS is not a member of this new family.

Our scheme is to take as our dual set the intersection of two unit balls in different \mathcal{L}^p norms, one ball at $\mathbf{0}$ and the other ball at $\mathbf{1}$. For $1 < p < \infty$, the intersection of the two balls is a strictly convex set that includes the unit axes but does not exceed the unit hypercube. (At $p = 1$, we get the probability simplex, which is not strictly convex. At $p = \infty$ we get the unit hypercube, which is also not strictly convex.) Let $\|\cdot\|_p$ denote the \mathcal{L}^p norm. Then we can define the vectors in the intersections of the unit balls, $\mathcal{U}(p)$, as

$$\mathcal{U}(p) \equiv \{\mathbf{y} \mid \mathbf{y} \in \mathbb{R}^n, \|\mathbf{y} - \mathbf{1}\|_p \leq 1, \|\mathbf{y}\|_p \leq 1\}$$

This set gives us a cost function $C(\mathbf{x}) = \max_{\mathbf{y} \in \mathcal{U}(p)} \mathbf{x} \cdot \mathbf{y}$. We dub this the *unit ball market maker*. Since we can easily test whether a vector is within both unit balls (i.e., within $\mathcal{U}(p)$), the optimization problem for the cost function can be solved in polynomial time.

This family of market makers is parameterized by the \mathcal{L}^p norm that defines which vectors in dual space are in the convex set. By choosing the value of p correctly, we can engineer a market maker with the desired maximum sum of prices. The outer boundary of the set is defined by the unit ball from $\mathbf{0}$ in \mathcal{L}^p space. Its boundary along $\mathbf{1}$ is given by the k that solves $\sqrt[p]{nk^p} = 1$. Solving for k we get $k = n^{-1/p}$, and so the maximum sum of prices is $nk = n(n^{-1/p}) = n^{1-1/p}$. For prices that are at most $1 + v$, we can set $1 + v = n^{1-1/p}$. Solving this equation for p yields $p = \frac{\log n}{\log n - \log(1+v)}$. Given any target maximum level of vigorish, this formula provides the exponent of the unit ball market maker to use. Considering that only small divergences away from unity are natural to the setting, the p we select for our \mathcal{L}^p norm should be quite small. The norm increases in the maximum sum of prices, and for larger n the same norm produces larger sums of prices.

One of the advantages of the unit ball market maker is that it is defined over all of \mathbb{R}^n , as opposed to just the non-negative orthant. Its behavior in the positive orthant is to charge agents more than a dollar for a dollar guaranteed payout, because the outer boundary diverges outwards from the probability simplex. Its behavior in the negative orthant, where its points on the inner boundary are selected in the maximization, is to pay less than a dollar for a dollar guaranteed payout. Its behavior in all other orthants is equivalent to max, as the unit axes are selected as maximizing arguments. Finally, if we restrict the unit ball market maker to only the non-negative orthant (like the OPRS), the sum of prices is tightly bounded between 1 and $n^{1-1/p}$.

6 Conclusions and future work

Using five desiderata that have appeared in the finance and prediction market literature, we contextualized a new class of cost functions, which we dubbed *homogeneous risk measures*. We showed that the OPRS [Othman et al., 2010] is a member of this class, because it is convex, monotonic, and positive homogeneous. We proved only the max cost function satisfies all five of our desiderata, but it does not have a differentiable price response. To produce a differentiable price response, one can add a regularizer, leading to the regularized online learning algorithms explored by Chen and Vaughan [2010]. Another approach is to curve the conjugate dual space, relaxing it from the probability simplex. We discussed how the properties of the convex set induce desirable properties in its conjugate homogeneous risk measure. Finally, using our insights, we developed a new family of homogeneous risk measures, the *unit ball market makers*, with desirable properties.

Our work centered on cost functions that are positive homogeneous, because these are the only cost functions that display identical relative price responses at different levels of liquidity. However, another direction is to explore cost functions that display some characteristics of liquidity sensitivity (more muted price responses at high levels of liquidity) without necessarily being homogeneous.

Finally, we are attracted to the work of Agrawal et al. [2009] because it provides a framework to simply add functionality to handle limit orders (orders of the form “I will pay no more than p for the payout vector \mathbf{x} ”) into a cost function market maker. That framework relies on convex optimization and so would also be able to run in polynomial time, a significant gain over naïve implementations of limit orders within cost function market makers. However, that work relied heavily on simplifications to the optimization that could be made because of translation invariance, so it is unclear how to embed a market maker whose convex conjugate is defined over more than the probability simplex into a limit order framework.

Acknowledgements We thank Peter Carr, Yiling Chen, Geoff Gordon, Dilip Madan, Dave Pennock, Steve Shreve, and Kevin Waugh for helpful discussions and suggestions. This work was supported by NSF grants CCF-1101668, IIS-0905390, IIS-0964579, and by the Google Fellowship in Market Algorithms.

Bibliography

- J. Abernethy, Y. Chen, and J. W. Vaughan. An optimization-based framework for automated market-making. In *ACM Conference on Electronic Commerce (EC)*, 2011.
- S. Agrawal, E. Delage, M. Peters, Z. Wang, and Y. Ye. A unified framework for dynamic pari-mutuel information market design. In *ACM Conference on Electronic Commerce (EC)*, pages 255–264, 2009.
- P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.

- A. Ben-Tal and M. Teboulle. An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476, 2007.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- P. Carr, H. Geman, and D. Madan. Pricing and hedging in incomplete markets. *Journal of Financial Economics*, 62(1):131–167, 2001.
- Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 49–56, 2007.
- Y. Chen and J. W. Vaughan. A new understanding of prediction markets via no-regret learning. In *ACM Conference on Electronic Commerce (EC)*, pages 189–198, 2010.
- Y. Chen, L. Fortnow, N. Lambert, D. M. Pennock, and J. Wortman. Complexity of combinatorial market makers. In *ACM Conference on Electronic Commerce (EC)*, pages 190–199, 2008.
- H. Föllmer and A. Schied. *Stochastic Finance (Studies in Mathematics 27)*. De Gruyter, 2002.
- S. Goel, D. Pennock, D. Reeves, and C. Yu. Yoopick: a combinatorial sports prediction market. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1880–1881, 2008.
- R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
- H. Markowitz. Portfolio Selection. *The Journal of Finance*, 7(1):77–91, 1952.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. In *ACM Conference on Electronic Commerce (EC)*, pages 253–254, 2009.
- A. Othman and T. Sandholm. Automated market-making in the large: the Gates Hillman prediction market. In *ACM Conference on Electronic Commerce (EC)*, pages 367–376, 2010a.
- A. Othman and T. Sandholm. When Do Markets with Simple Agents Fail? In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 865–872, Toronto, Canada, 2010b.
- A. Othman, D. M. Pennock, D. M. Reeves, and T. Sandholm. A practical liquidity-sensitive automated market maker. In *ACM Conference on Electronic Commerce (EC)*, pages 377–386, 2010.
- D. Pennock and R. Sami. Computational Aspects of Prediction Markets. In *Algorithmic Game Theory*, chapter 26, pages 651–674. Cambridge University Press, 2007.
- M. Peters, A. M.-C. So, and Y. Ye. Pari-mutuel markets: Mechanisms and performance. In *International Workshop On Internet And Network Economics (WINE)*, pages 82–95, 2007.
- R. T. Rockafellar. Level sets and continuity of conjugate convex functions. *Transactions of the American Mathematical Society*, 123(1):46–63, 1966.
- R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- S. Shalev-Shwartz and Y. Singer. A primal-dual perspective of online learning algorithms. *Machine Learning*, 69:115–142, December 2007.