Proceedings Article

An Axiomatic Characterization of Adaptive-Liquidity Market Makers

XIAOLONG LI, University of Texas at Austin
JENNIFER WORTMAN VAUGHAN, Microsoft Research, New York City & UCLA

Prediction markets offer contingent securities with payoffs linked to future events. The market price of a security reveals information about traders' beliefs, but prediction markets often suffer from low liquidity, which can prevent trades from occurring. One way to inject liquidity into a market is through the use of an automated market maker, an algorithmic agent willing to accept some risk in order to facilitate trades. Abernethy et al. [2013] proposed a general framework for the design of automated market makers, defining a set of axioms that a market should satisfy, and characterizing the class of market makers that satisfy these axioms. However, the liquidity of any market in their class, quantified in terms of the rate at which prices adapt to trades, is fixed a priori and does not change as the volume of trade increases. Othman and Sandholm [2011] proposed a class of liquidity-adaptive markets, but gave little guidance for how to choose a market from within this class. Combining ideas from Abernethy et al. [2013] and Othman and Sandholm [2011], we provide an axiomatic characterization of a parameterized class of automated market makers with adaptive liquidity. A primary advantage of our framework is the ability to analyze important market properties, such as its ability to aggregate information or make a profit, in terms of market parameters, which we do using techniques from convex analysis and geometry. For example, we show that the curvature of the price space can be used to manage a trade-off between information loss and profit. This analysis offers guidance for market designers who wish to choose a particular market maker to implement.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Automated market makers, prediction markets, convex analysis

1. INTRODUCTION

A securities market offers a set of *contingent securities* with payments associated with future events. For example, an Oscars market may offer a security that is worth \$1 if *Argo* wins the Academy Award for Best Picture and \$0 otherwise. A risk neutral trader who believes that the probability of *Argo* being named Best Picture is p would value such a security at \$$p$. The price at which securities are traded therefore reveals information about the traders' beliefs. When emphasizing the aggregation of traders' beliefs, securities markets are sometimes referred to as *prediction markets*.

Security prices can be determined in a variety of ways. The popular prediction market website Intrade prices securities through a continuous double auction, similar to the stock market; traders submit limit orders, and buyers are matched with sellers in real time. This approach is reasonable in thick markets, but can be problematic when the number of traders is relatively small. Traders may be unwilling to wait around in the market until a suitable trading partner arrives. Once this starts causing traders to leave the market, the problem is exacerbated [Pennock and Sami 2007].

This research was partly completed while Li was at UCLA and partially supported by the NSF under grant IIS-1054911. Any opinions, findings, conclusions, or recommendations are those of the authors alone.

EC'13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

One way to get around this problem is to operate the market using an automated market maker. An automated market maker is an algorithmic agent that adaptively sets prices for each security based on the history of trades, and is always willing buy or sell securities at those prices. By taking on some risk, the market maker effectively injects liquidity into the market, which encourages trader participation.

Hanson [2003, 2007] introduced *market scoring rules*, a class of automated market makers for *complete markets*, in which there is one security associated with each possible state of the world. Market scoring rules can be viewed as a generalization of proper scoring rules to multiple forecasters. Abernethy et al. [2011, 2013] proposed a more general framework for the design of market makers for securities markets.[1] They began by defining a set of axioms formalizing intuitive properties that a "reasonable" market should satisfy, and showed that any market maker satisfying these axioms must price securities via a convex potential function of a certain form, referred to as the cost function. Using conjugate duality, they showed how to express any cost function of this form as an optimization over the space of feasible prices. This is described in more detail in Section 2.

These market makers inject liquidity into the market as desired. However, they share a property referred to by Othman and Sandholm [2011] as *liquidity insensitivity*: the liquidity of the market (quantified in terms of the rate at which prices adapt to trades) is fixed a priori and does not change as the volume of trade increases. In fact, the prices set by the market maker depend only on the *differences* between the quantities of various securities that have been purchased. As an extreme example, consider a market offering one security (D) that is worth \$1 if and only if a Democrat is elected in the 2016 US Presidential election, and a second (R) worth \$1 if and only if a Republican is elected. A trader who wished to purchase security D after 100 units D and 0 of R had been purchased would pay the same price as a trader who entered the market after 10,100 units of D and 10,000 of R had been purchased, and his trade would have the same impact on future market prices. Market liquidity impacts traders' ability to move market prices as well as the market maker's risk, and in practice, choosing the right amount of liquidity a priori is viewed as a black art [Pennock 2010].

As an alternative, Othman et al. [2010] introduced a *liquidity-sensitive* market maker based on Hanson's logarithmic market scoring rule. The key property of this market maker is that price movement slows as the volume of trade grows. Othman and Sandholm [2011] expanded on this idea to propose a broader class of liquidity-sensitive markets which they call *homogeneous risk measures*. However, little guidance is given for choosing the best market from within the proposed class.

Combining ideas from the duality-based market makers of Abernethy et al. [2013] and the homogeneous risk measures of Othman and Sandholm [2011], we introduce a generalized parameterized class of automated market makers with adaptive liquidity. Following Abernethy et al., we take an axiomatic approach and derive a characterization of liquidity adaptive markets. The markets in our framework behave like duality-based markets initially, but like homogeneous risk measures in the limit as the volume of trade grows large. Thus we get the benefits of both; early traders have more incentive to participate in our markets compared with homogeneous risk measures, but price movement slows over time.

Aside from characterizing the parameterized class of liquidity-adaptive market makers satisfying our axioms, the primary contribution of our research is an in-depth analysis of the market makers in this class. This analysis offers guidance for designers who wish to choose a particular market maker to implement. We look at two factors:

- First, we outline methods of quantifying the extent to which a particular market can aggregate information about traders' beliefs. We show that any liquidity-adaptive market must sacrifice *expressiveness*, a property of duality-based markets that allows traders to

[1]In the complete market setting, their markets are equivalent to market scoring rules in a particular strong sense [Abernethy et al. 2013], but their framework also includes incomplete markets.

"push" market prices to match their beliefs. Without expressiveness, there can be a range of beliefs (i.e., distributions over outcomes) consistent with current market prices in the sense that risk-neutral traders who have those beliefs have no incentive to trade. The size of this set provides a measure of how much information is lost. For different notions of "size," which relate to the goal of the market designer, we show how the parameters of our market impact this information loss.
- Second, we investigate profitability. We describe and quantify three necessary factors for the market maker to make a profit: sufficient trade volume, disagreement among traders, and *curvature of the price space*. The market designer can use curvature as a knob to adjust the trade-off between the market's potential profitability and its information loss.

We note that our measure of information loss can be viewed as softening of expressiveness. When it's zero, absolute expressive is achieved. When it's non-zero but bounded, the market still provides useful information but liquidity adaptivity and profitability become achievable.

The proofs of these results are of a different flavor than previous analyses of the markets mentioned above. In particular, many of our proofs rely on geometric intuitions, examining the role of the shape of the market's price space.

Related to our work is the recent research of Othman and Sandholm [2012], which also examines the problem of designing liquidity-adaptive markets. They take an approach orthogonal to ours, beginning with a specific class of markets rather than providing an axiomatic characterization. They achieve liquidity adaptivity, but other properties of our markets, such as information incorporation or the existence and continuity of instantaneous prices cannot be established easily in their framework. They do not discuss the extent to which their markets are capable of aggregating information about traders' beliefs or how this interacts with profitability.

In Section 2, we review duality-based markets and homogeneous risk measures. In Section 3, we define a set of axioms, including liquidity adaptation, and give a formal characterization of a class of automated market makers that satisfy these axioms. In Section 4, we quantify the extent to which a market is capable of aggregating trader beliefs, and show how this relates to the parameters of our markets. We do the same for profit in Section 5, and provide guidance on how to balance profit with loss of information when choosing a market maker to implement. We conclude with remarks on how to extend this framework.

2. BACKGROUND AND RELATED WORK

We consider the design of market makers for complete markets over finite outcome spaces. In this setting, there are n mutually exclusive and exhaustive states of the world, indexed $1, \ldots, n$, and n securities, with each security i worth \$1 if outcome i occurs and \$0 otherwise. The automated market maker is always willing to sell or buy securities at prices that adapt over time. We begin by reviewing two classes of market makers for this setting, duality-based markets and homogeneous risk measures.

2.1. Duality-Based Markets and Cost Functions

Abernethy et al. [2011, 2013] proposed a set of five axioms that an automated market maker should satisfy: path independence, existence of instantaneous prices, information incorporation, no-arbitrage, and expressiveness. Loosely speaking, path independence says that a trader should pay the same price for a bundle of securities regardless of whether he purchases the full bundle at once or splits up his order into multiple consecutive purchases. Existence of instantaneous prices requires that the price per security of purchasing an infinitesimally small fraction of a security is well-defined; this price will correspond to an estimate of traders' beliefs. Information incorporation says that the cost of a security should increase as traders purchase that security. No-arbitrage says that it should never be possible for a trader to purchase a bundle of securities that allows him to make a positive profit regardless of the outcome. Finally, expressiveness says that a trader with enough wealth at his disposal should always be able to "push" the market prices to match his beliefs by buying and selling securities.

To characterize markets satisfying these axioms, Abernethy et al. first show that a market is path independent if and only if prices are determined by a potential function C, referred to as the *cost function*. Let $\mathbf{q}$ be a vector of length n with each component q_i indicating the number of securities associated with outcome i that have been purchased by traders in the market so far. (Here q_i can be negative if more units of the ith security have been sold than purchased, which can happen if short selling is allowed, and need not be an integer if fractional purchases are allowed.) The vector $\mathbf{q}$ summarizes the current market state. Now, suppose a trader would like to purchase a bundle of securities $\mathbf{r}$, with r_i denoting the number of securities associated with outcome i that he would like to purchase. (Here r_i could be 0, or in some cases, negative, representing a (short) sale, and need not be an integer value.) In a cost function based market, the amount that the trader must pay to the market maker for this bundle is $C(\mathbf{q}+\mathbf{r}) - C(\mathbf{q})$. Instantaneous prices exist if $\nabla C(\mathbf{q})$ is well-defined; the instantaneous price of security i, denoted $p_i(\mathbf{q})$, is the ith component of $\nabla C(\mathbf{q})$.

Let Δ_n denote the probability simplex. Abernethy et al. go on to show that, in the complete market setting, the remaining three axioms are satisfied if and only if it is possible to write the cost function C as $C(\mathbf{q}) = \sup_{\mathbf{p}\in\Delta_n} \mathbf{p}\cdot\mathbf{q} - R(\mathbf{p})$ for some strictly convex function R known in convex analysis as the *conjugate* of C. We refer to such markets as *duality-based* since they build on conjugate duality. Any cost function of this form is itself convex.

The well-studied *logarithmic market scoring rule* (LMSR) [Hanson 2003, 2007] is a special case of this framework, obtained by setting R to be the negative entropy function scaled by a parameter $b > 0$. Its cost function can be written as $C(\mathbf{q}) = b \log\left(\sum_{i=1}^{n} e^{q_i/b}\right)$, and prices are given by $p_i(\mathbf{q}) = e^{q_i/b}/\sum_{j=1}^{n} e^{q_j/b}$. The parameter b controls both the rate at which prices change and the worst-case monetary loss of the market maker.

It is easy to see that the sum over all securities i of the instantaneous prices $p_i(\mathbf{q})$ of the LMSR is always one. This property, sometimes referred to as *translation invariance* [Agrawal et al. 2011; Chen and Wortman Vaughan 2010], follows directly from the no-arbitrage condition, and therefore holds of any duality-based market maker in the framework of Abernethy et al. [2013]. Aside from preventing arbitrage, it allows us to interpret the vector of market prices directly as a probability distribution over outcomes; indeed, in such a market, a myopic risk neutral trader will have incentive to trade unless the prices match his beliefs, which allows information to be aggregated.

2.2. Homogeneous Risk Measures

One disadvantage of translation invariant market makers is that the liquidity of the market (measured in terms of the rate at which prices change as purchases are made) must be fixed in advance, before the activity level of traders is known. If prices react too quickly, the market will be unstable with many fluctuations when the number of traders is large. If prices react too slowly, traders may not have enough wealth to move the market prices to reflect their collective beliefs.

To address this problem, Othman et al. [2010] proposed the OSPR market maker, a slight modification of Hanson's LMSR in which the liquidity parameter b varies over time. The OSPR cost function is given by $C(\mathbf{q}) = \alpha b(\mathbf{q}) \log\left(\sum_{i=1}^{n} e^{q_i/b(\mathbf{q})}\right)$, where $b(\mathbf{q}) = \sum_{i=1}^{n} q_i$ and $\alpha > 0$ is a tunable parameter. The OSPR does not satisfy translation invariance. Instead, the sum of prices is *at least* one. Arbitrage can still be prevented by disallowing short selling. If only purchases are allowed, the difference between the sum of prices and one can be viewed as the "profit cut" of the market maker, which opens up the potential of profitability.

Building on this intuition, Othman and Sandholm [2011] proposed an alternative class of market makers, *homogeneous risk measures*, which are implemented using cost functions that can be written as $C(\mathbf{q}) = \sup_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot(\mathbf{q}+\mathbf{q}_0) = \sup_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot\mathbf{q}+\mathbf{p}\cdot\mathbf{q}_0$ for some $\mathbf{q}_0 \in \mathbb{R}^n_+ \setminus \mathbf{0}$, and some $\mathbb{Y} \subset \mathbb{R}^n_+$ with a strictly convex boundary.[2] The cost function of a homogeneous risk measure is, as the name suggests, positive homogeneous: for all $\lambda \geq 0$, $C(\lambda\mathbf{q}) = \lambda C(\mathbf{q})$. One way to measure liquidity is by the speed at which the price vector $\mathbf{p}$ changes as the quantity vector $\mathbf{q}$ changes, which can be captured by $(\partial p_i/\partial q_i)^{-1}$. Homogeneity implies that this measure of liquidity increases linearly with the size of quantity vector.

2.3. A Broader Class of Adaptive-Liquidity Markets

A major goal of this research is to characterize the class of markets that satisfy the nice properties of duality-based markets, yet simultaneously provide adaptive liquidity like homogeneous risk measures. The conjugate function R that is used to define the cost function of a duality-based market adds stability to market prices, preventing rapid fluctuation as trades are made.[3] A similar role is played by the curved space $\mathbb{Y}$ in the cost function of a homogeneous risk measure. (The linear term $\mathbf{p}\cdot\mathbf{q}_0$ does not add stability, and is necessary only to make the function differentiable at $\mathbf{q} = \mathbf{0}$.) The cost functions we derive in Section 3 incorporate both simultaneously. In particular, we will arrive at cost functions of the form $C(\mathbf{q}) = \sup_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot\mathbf{q} - R(\mathbf{p})$. The effect of R dominates when $\|\mathbf{q}\|$ is small, whereas the effect of $\mathbb{Y}$ dominates when $\|\mathbf{q}\|$ is large. We will see that under certain restrictions on $\mathbb{Y}$ and R, cost functions of this form behave asymptotically like homogeneous risk measures and can thus provide increasing liquidity as the size of $\mathbf{q}$ grows.

3. AN AXIOMATIC CHARACTERIZATION

In this section, we characterize the class of automated market makers that satisfy the first four axioms of Abernethy et al. [2013] (path independence, existence of instantaneous prices, information incorporation, and no-arbitrage), have bounded worst-case loss for the market maker, and have adaptive liquidity. Note that we do not guarantee expressiveness; the inherent conflict between liquidity adaptivity and expressiveness and the implications for information aggregation are discussed in Section 4.

3.1. The Buy-Only Principle

As discussed above, in order to allow liquidity to adapt over time, Othman et al. [2010] and Othman and Sandholm [2011] allow security prices to sum to more than one. In order to prevent arbitrage, it is then necessary to disallow the sale of securities; if selling were permitted, traders could sell one each of every security and be guaranteed a profit. We adapt the convention of disallowing sales as well.

Disallowing sales has another benefit that we exploit. As described in Section 2.1, in a path independent (i.e., cost function based) market, the current market state is summarized by the vector $\mathbf{q}$ of outstanding shares of each security. If both purchases and sales are allowed, then $||\mathbf{q}||$ could be small for two reasons: either there has not been much trade in the market (i.e., the market is thin), or traders disagree, resulting in similar volumes of purchases and sales (in which case the market is thick). Since we cannot distinguish these two cases based on $\mathbf{q}$ alone, the market must behave similarly in both cases. On the other hand, if traders are restricted to purchases only, then a small value of $||\mathbf{q}||$ unambiguously implies a thin market. More generally, $||\mathbf{q}||$ can be used to track the volume of trade in the market, allowing us to adapt liquidity directly based on trade volume.

This restriction is not as limiting as it may seem. Suppose a trader would like to sell the security associated with outcome i. If he instead purchases one of every security $j \neq i$ for

[2] Throughout this paper, we use $\mathbb{R}^n_+$ to denote $\{\mathbf{x} \in \mathbb{R}^n \mid x_i \geq 0,\ \forall i\}$, and $\mathbb{R}^n_{++}$ to denote the subset of $\mathbb{R}^n_+$ in which no coordinate is equal to 0.

[3] Specifically, R plays a mathematical role identical to that of the *regularizers* frequently employed by the machine learning and statistics communities [Chen and Wortman Vaughan 2010; Abernethy et al. 2013].

some price c, then his total payoff will be $-c$ if outcome i occurs and $1 - c$ if any outcome $j \neq i$ occurs. This is equivalent to selling security i at price $1-c$. However, unlike a standard duality-based market, the selling price will typically be strictly lower than the purchasing price, leading to a non-zero bid-ask spread.

For these reasons, we consider markets in which selling is disallowed from this point forward. We refer to such markets as *buy-only markets* when there may be confusion.

3.2. Adapting the Axioms

We first state formally the axioms from Section 2.1. These axioms were originally defined for markets with short selling, but can easily be adapted to buy-only markets as we do here.

Path Independence and the Use of Cost Functions:

We begin with path independence, the property that a trader pays the same price to purchase a bundle r regardless of how he splits up the purchase.

PROPERTY 1 (PATH INDEPENDENCE). $\forall \mathbf{r}, \mathbf{r}', \mathbf{r}'' \in \mathbb{R}_+^n$ *s.t.* $\mathbf{r} = \mathbf{r}' + \mathbf{r}''$, *given any fixed history of previous market transactions, the cost to a trader of purchasing bundle* $\mathbf{r}$ *is the same as the cost of purchasing* $\mathbf{r}'$ *and then immediately purchasing* $\mathbf{r}''$.

Abernethy et al. [2013] showed that path independence holds if and only if costs can be calculated via a potential function C, which remains true for buy-only markets.

PROPOSITION 3.1. *In buy-only markets, Path Independence holds if and only if there exists a function* $C : \mathbb{R}^n \to \mathbb{R}$ *such that cost of purchasing a bundle of securities* $\mathbf{r} \in \mathbb{R}_+^n$ *given that traders have previously purchased bundles* $\mathbf{r}_1, \cdots, \mathbf{r}_T$ *is* $C(\sum_{t=1}^T \mathbf{r}_t + \mathbf{r}) - C(\sum_{t=1}^T \mathbf{r}_t)$.

The proof is identical to the proof of Theorem 3.1 of Abernethy et al. [2013]. As in previous work, we restrict our attention to cost function based markets, and assume a cost function C exists when defining additional properties.

Information Incorporation and the Use of Conjugate Duality:

Information incorporation says that the purchase of a bundle $\mathbf{r}$ should never cause the price of $\mathbf{r}$ to decrease.

PROPERTY 2 (INFORMATION INCORPORATION). $\forall \mathbf{q}, \mathbf{r} \in \mathbb{R}_+^n$, $C(\mathbf{q} + 2\mathbf{r}) - C(\mathbf{q} + \mathbf{r}) \geq C(\mathbf{q} + \mathbf{r}) - C(\mathbf{q})$.

Abernethy et al. [2013] showed that information incorporation holds if and only if C is convex. This also remains true in buy-only markets. The proof is identical to the start of the proof of Theorem 3.2 in Abernethy et al. [2013].

PROPOSITION 3.2. *A buy-only market implemented via a cost function* $C : \mathbb{R}_+^n \to \mathbb{R}$ *satisfies Information Incorporation if and only if* C *is convex.*

The convexity of the cost function C allows us to make use of powerful tools from convex analysis. To do so, it will be convenient to restrict our attention to closed cost functions. Any function has a closed function that is equivalent to it everywhere except possibly at the relative boundary of its domain, so we do not lose any power by assuming that C is closed. Under this assumption, since C must be convex, it follows from basic results of convex analysis that we can write C as $C(\mathbf{q}) = \sup_{\mathbf{p} \in \mathbb{Y}} \mathbf{p} \cdot \mathbf{q} - R(\mathbf{p})$ for some closed and convex function R and convex set $\mathbb{Y} \subseteq \mathbb{R}^n$ [Rockafellar 1997]. R is called the conjugate of C. We call such markets *duality-based buy-only markets*.

Definition 3.3 (*Duality-Based Buy-Only Market*). A *duality-based buy-only market* is a buy-only market implemented via a closed, convex cost function C of the form

$$C(\mathbf{q}) = \sup_{\mathbf{p} \in \mathbb{Y}} \mathbf{p} \cdot \mathbf{q} - R(\mathbf{p}) \tag{1}$$

for some convex set $\mathbb{Y} \subseteq \mathbb{R}^n$ and some closed, convex function $R : \mathbb{Y} \to \mathbb{R}$.

In some of the results that follow, we assume that R is bounded. Since R is already closed, $\mathbb{Y} = \mathrm{dom}(R)$ is also closed by standard convex analysis arguments. Since R is bounded, $\mathbb{Y}$ must be bounded along all directions; otherwise the optimization problem for C would yield $+\infty$. Bounded along all directions is the same as bounded for a convex set, so $\mathbb{Y}$ must be closed and bounded, and hence compact. Whenever we assume the boundedness of R in a result, it should be understood that $\mathbb{Y}$ is then compact.

The remaining axioms will impose additional constraints on the choice of R and $\mathbb{Y}$.

Sufficient Conditions for Existence of Instantaneous Prices:

As we will discuss in detail in Section 4, the market price of a security provides information about traders' beliefs about the probability that the associated outcome will occur. Since the goal of running a prediction market is to aggregate beliefs, it is important that the notion of market prices be well-defined. Since prices change as trades are made in a cost function based market, we require well-defined *instantaneous* prices. The instantaneous price of security i is the per unit price of purchasing an infinitesimally small quantity of the security, i.e., the derivative of C with respect to q_i.

PROPERTY 3 (EXISTENCE OF INSTANTANEOUS PRICES). *C is continuous and differentiable everywhere on $\mathbb{R}^n_+$.*

It is known in convex analysis that a function of the form of Equation 1 is differentiable at some point $\mathbf{q}$ if and only if the maximization has a unique optimal solution on $\mathbb{Y}$ [Rockafellar 1997]. In this case, the instantaneous price vector, denoted $\mathbf{p}(\mathbf{q})$ is the unique maximizer of the optimization, i.e.,

$$\mathbf{p}(\mathbf{q}) := \nabla C(\mathbf{q}) = \arg\max_{\mathbf{p} \in \mathbb{Y}} \mathbf{p} \cdot \mathbf{q} - R(\mathbf{p}). \tag{2}$$

Thus a cost function C is differentiable on $\mathbb{R}^n_+$ if and only if the maximization has a unique optimal solution on $\mathbb{Y}$ for every $\mathbf{q} \in \mathbb{R}^n_+$. Technically, one could show that instantaneous prices exist if and only if R is strictly convex at any point $\mathbf{x}$ such that the subgradient at $\mathbf{x}$ intersects $\mathbb{R}^n_+$. This characterization, though accurate, is awkward to use. For the purpose of market design, it will be more useful to instead provide two *sufficient* conditions for instantaneous prices to exist, given in Propositions 3.4 and 3.5.

To state these conditions, it is useful to introduce some notation. Let $\partial\mathbb{Y}$ denote the boundary of the set $\mathbb{Y}$. We define the *positive boundary* of $\mathbb{Y}$, denoted $\partial^+\mathbb{Y}$, as

$$\partial^+\mathbb{Y} := \{\mathbf{y} \in \mathrm{cl}\mathbb{Y} \mid \nexists \mathbf{x} \in \mathbb{Y} \text{ } s.t. \text{ } \mathbf{x} \neq \mathbf{y} \text{ } and \text{ } \mathbf{x} \succeq \mathbf{y}\}.$$

The positive boundary of a set consists of those points that are not *dominated* by any other point in the set, where $\mathbf{x}$ dominates $\mathbf{y}$ if and only if $\mathbf{x} \succeq \mathbf{y}$. Clearly $\partial^+\mathbb{Y} \subseteq \partial\mathbb{Y}$. Finally, we say that a set is *strictly non-linear* if it contains no line segment.

Proposition 3.4 deals with the special case of homogeneous risk measures.

PROPOSITION 3.4. *A duality-based buy-only market with $R(\mathbf{p}) = \mathbf{p} \cdot \mathbf{q}_0$ for some $\mathbf{q}_0 \in \mathbb{R}^n_{++}$ satisfies Existence of Instantaneous Prices if $\partial^+\mathbb{Y}$ is strictly non-linear.*

PROOF. For any $\mathbf{q} \in \mathbb{R}^n_+$, we have $\mathbf{q} + \mathbf{q}_0 \in \mathbb{R}^n_{++}$. Suppose a point $\mathbf{x}$ is the maximizer of $\mathbf{p} \cdot (\mathbf{q} + \mathbf{q}_0)$ on $\mathbb{Y}$. Then $\mathbf{x}$ cannot be dominated by another point $\mathbf{y} \in \mathbb{Y}$; otherwise $\mathbf{y}$ would yield a larger objective value. Thus $\mathbf{x}$ must be in $\partial^+\mathbb{Y}$ by definition. For a concave objective function, the maximizers must form a convex set. If $\mathbf{x}_1$ and $\mathbf{x}_2$ are distinct maximizers, then any point on the line segment connecting $\mathbf{x}_1$ and $\mathbf{x}_2$ must be a maximizer and therefore on $\partial^+\mathbb{Y}$, contradicting our assumption that $\partial^+\mathbb{Y}$ is strictly non-linear. Therefore, the maximizer must be unique for all $\mathbf{q} \in \mathbb{R}^n_+$. $\square$

Proposition 3.5 moves beyond homogeneous risk measures.

Proposition 3.5. *Any duality-based buy-only market with strictly convex R satisfies Existence of Instantaneous Prices.*

Proof. The strict convexity of R implies that the objective function as a whole is strictly concave, which immediately guarantees the uniqueness of maximizer. $\square$

Bounded Worst-Case Loss:

The next axiom we consider is bounded worst-case loss of the market maker. This is not stated formally as an axiom by Abernethy et al. [2013], but holds in their framework as long as R is bounded since $\mathbb{Y}$ is restricted to be the probability simplex.

Property 4 (Bounded Loss). *The worst-case loss of the market maker, $\sup_{\mathbf{q}\in\mathbb{R}^n_+}(\max_{i\in\{1,\cdots,n\}} q_i - C(\mathbf{q}) + C(\mathbf{0}))$, is bounded.*

Define the *extension* of a set $\mathbb{Y}$, denoted $extend(\mathbb{Y})$, to be the set of all points that are dominated by a point in $\mathbb{Y}$. Formally, $extend(\mathbb{Y}) = \{\mathbf{x} \in \mathbb{R}^n \mid \exists \mathbf{y} \in \mathbb{Y},\ \mathbf{y} \succeq \mathbf{x}\}$. In buy-only markets, when R is bounded, it is possible to obtain bounded worst-case loss even if the probability simplex is not contained in $\mathbb{Y}$ as long as it is contained in its extension.

Proposition 3.6. *A duality-based buy-only market with a bounded R satisfies Bounded Loss if and only if $\Delta_n \subseteq extend(\mathbb{Y})$. In particular, if $\Delta_n \subseteq extend(\mathbb{Y})$, then the loss is bounded by $\sup_{p\in\mathbb{Y}} R(p) - \inf_{p\in\mathbb{Y}} R(p)$.*

Proof. First we will show that if $\Delta_n \subseteq extend(\mathbb{Y})$, then the market maker has bounded loss. Suppose that $\Delta_n \subseteq extend(\mathbb{Y})$. Then for each i, there exists some $\mathbf{p} \in \mathbb{Y}$ such that $\mathbf{e}_i \preceq \mathbf{p}$. By Equation 1, for any $\mathbf{q}$, $C(\mathbf{q}) = \sup_{\mathbf{p}'\in\mathbb{Y}} \mathbf{p}' \cdot \mathbf{q} - R(\mathbf{p}') \geq \mathbf{p} \cdot \mathbf{q} - R(\mathbf{p}) \geq \mathbf{e}_i \cdot \mathbf{q} - R(\mathbf{p}) = q_i - R(\mathbf{p}) \geq q_i - M$, where $M = \sup_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p})$, which must be finite since we have assumed that R is bounded. We therefore have that that $C(\mathbf{q}) \geq q_i - M$ for any $\mathbf{q}$ and any i. Similarly, letting $m = \inf_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p})$, we have $C(\mathbf{0}) = \sup_{\mathbf{p}\in\mathbb{Y}} -R(\mathbf{p}) = -\inf_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p}) = -m$, which also must be finite. Therefore, $\sup_{\mathbf{q}\in\mathbb{R}^n_+}(\sup_{i\in\{1,\cdots,n\}} q_i - C(\mathbf{q}) + C(\mathbf{0})) \leq M - m < \infty$.

We prove the other direction by contradiction. Suppose there exists some $\mathbf{x} \in \Delta_n$ such that $\mathbf{x} \notin extend(\mathbb{Y})$. Notice that $extend(\mathbb{Y})$ is convex. Since we have assumed that $\mathbb{Y}$ is compact, $extend(\mathbb{Y})$ is also closed. Since $\mathbf{x}$ is also closed and $\mathbf{x} \notin extend(\mathbb{Y})$, there must exist a hyperplane $P = \{\mathbf{y} \mid \mathbf{y} \cdot \mathbf{b} = \beta\}$ for some $\mathbf{b}$ and β that strongly separates $\mathbf{x}$ and extend($\mathbb{Y}$), so that $\mathbf{x} \cdot \mathbf{b} > \beta > \max_{\mathbf{p}\in\mathbb{Y}} \mathbf{p} \cdot \mathbf{b}$ [Rockafellar 1997]. Notice that $\mathbf{b}$ must belong to $\mathbb{R}^n_+ \setminus \{\mathbf{0}\}$, otherwise if $b_i < 0$, we can find some $\mathbf{p} \in \text{extend}(\mathbb{Y})$ where p_i is negative with large absolute value such that $\mathbf{p} \cdot \mathbf{b} > \beta$. Let $\delta = \mathbf{x} \cdot \mathbf{b} - \max_{\mathbf{p}\in\text{extend}(\mathbb{Y})} \mathbf{p} \cdot \mathbf{b} > 0$. Then for any $k \in \mathbb{R}$, $C(k\mathbf{b}) = \max_{\mathbf{p}\in\mathbb{Y}} k(\mathbf{p} \cdot \mathbf{b}) - R(\mathbf{p}) \leq k(\mathbf{x} \cdot \mathbf{b} - \delta) + \sup_{\mathbf{p}\in\mathbb{Y}} |R(\mathbf{p})| \leq k(\mathbf{x} \cdot (\max_i b_i)\mathbf{1}) - k\delta + \sup_{\mathbf{p}\in\mathbb{Y}} |R(\mathbf{p})| = k \max_{i\in\{1,\cdots,n\}} b_i - k\delta + \sup_{\mathbf{p}\in\mathbb{Y}} |R(\mathbf{p})|$. We can then lower bound the worst-case loss as

$$\sup_{\mathbf{q}\in\mathbb{R}^n_+}\left(\max_{i\in\{1,\cdots,n\}} q_i - C(\mathbf{q}) + C(\mathbf{0})\right) \geq \sup_{\mathbf{q}=k\mathbf{b},k\in\mathbb{R}^+}\left(\max_{i\in\{1,\cdots,n\}} q_i - C(\mathbf{q}) + C(\mathbf{0})\right)$$
$$= \sup_{k\in\mathbb{R}^+}(k \max_{i\in\{1,\cdots,n\}} b_i - C(k\mathbf{b}) + C(\mathbf{0})) \geq \sup_{k\in\mathbb{R}^+} k\delta - \sup_{\mathbf{p}\in\mathbb{Y}} |R(\mathbf{p})| - \inf R(\mathbf{p}),$$

which is unbounded since $\sup_{k\in\mathbb{R}^+} k\delta$ is unbounded but R is not. $\square$

No Arbitrage:

Finally, it is desirable to prevent traders from exploiting the market to make guaranteed profits. The no-arbitrage axiom of Abernethy et al. [2013] can easily be extended to buy-only markets. This property guarantees that a trader can never purchase a bundle that guarantees a positive profit on every possible outcome.

PROPERTY 5 (NO ARBITRAGE). $\forall \mathbf{q}, \mathbf{r} \in \mathbb{R}_+^n$, $\exists i \in \{1, \cdots, n\}$ *s.t.* $C(\mathbf{q}+\mathbf{r}) - C(\mathbf{q}) \geq r_i$.

The following proposition gives sufficient and necessary conditions for the market to satisfy No Arbitrage assuming that the cost function is differentiable. To simplify presentation, let $S_n = \{x \in \mathbb{R}_+^n \mid \sum_{i=1}^n x_i < 1\}$. We will refer to S_n as the *sub-probability simplex*. The result says that $\mathbb{Y}$ must be convex and in the positive orthant but outside the sub-probability simplex, or that we can replace $\mathbb{Y}$ with another set $\mathbb{Y}'$ satisfying these properties without changing the behavior of the cost function. Intuitively, the instantaneous prices must never sum to less than 1.

PROPOSITION 3.7. *Any duality-based buy-only market that satisfies Existence of Instantaneous Prices satisfies No Arbitrage if and only if there exists a convex set* $\mathbb{Y}' \subseteq \mathbb{R}_+^n \setminus S_n$ *such that for all* $\mathbf{q}$, $C(\mathbf{q}) = \sup_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot\mathbf{q} - R(\mathbf{p}) = \sup_{\mathbf{p}\in\mathbb{Y}'} \mathbf{p}\cdot\mathbf{q} - R(\mathbf{p})$.

PROOF. We first show that No Arbitrage implies that there exists a $\mathbb{Y}'$ satisfying the properties in the proposition statement.

Suppose that the market satisfies No Arbitrage. From simple analysis, we know that if a convex function is differentiable, it must be continuously differentiable, and so the price mapping $\mathbf{q} \to \mathbf{p}(\mathbf{q})$ is continuous. Suppose there exists some $\hat{\mathbf{q}}$ such that the ith component of $\mathbf{p}(\hat{\mathbf{q}})$ is negative. Then by continuity of price, there exists some small purchase $\delta\mathbf{e}_i$ such that the ith component of $\mathbf{p}(\hat{\mathbf{q}}+\theta\mathbf{e}_i)$ remains negative for all $\theta \in [0, \delta]$. The purchase of $\delta\mathbf{e}_i$ at market state $\hat{\mathbf{q}}$ is an arbitrage opportunity; since the cost of the trade is the integral of the price function, the trader receives money for the trade, and never has to pay money after the outcome is announced. Therefore, any $\mathbf{p} \notin \mathbb{R}_+^n$ can never be the price, or by Equation 2, the maximizer in the optimization, and so replacing $\mathbb{Y}$ with $\mathbb{Y} \cap \mathbb{R}_+^n$ will not change C.

Similarly, if $\exists \bar{\mathbf{q}} \in \mathbb{R}_+^n$, such that $\mathbf{p}(\bar{\mathbf{q}}) \in S_n$, then $\mathbf{p}(\bar{\mathbf{q}}) \cdot \mathbf{1} < 1$. By continuity of price, $\exists \delta > 0$ such that $\mathbf{p}(\bar{\mathbf{q}}+\theta\mathbf{1})\cdot\mathbf{1} < 1$ for $\theta \in [0, \delta]$. Let $f(t) = C(\bar{\mathbf{q}}+t\mathbf{1})$, so $f'(t) = \mathbf{p}(\bar{\mathbf{q}}+t\mathbf{1})\cdot\mathbf{1}$. Then $C(\bar{\mathbf{q}}+\delta\mathbf{1}) - C(\bar{\mathbf{q}}) = f(\delta) - f(0) = \int_0^\delta \mathbf{p}(\bar{\mathbf{q}}+\theta\mathbf{1})\cdot\mathbf{1}\mathrm{d}\theta < \int_0^\delta 1\mathrm{d}\theta = \delta$. This implies that the purchase of $\delta\mathbf{1}$ at market state $\bar{\mathbf{q}}$ is an arbitrage opportunity, since the trader pays less than δ but is guaranteed a payout of δ regardless of the outcome. Therefore, we can replace $\mathbb{Y}$ with $\mathbb{Y} \cap (\mathbb{R}_+^n \setminus S_n)$ without changing C. The convexity of $\mathbb{Y}$ is kept since $\mathbb{R}_+^n \setminus S_n$ is convex and intersection of convex sets is still convex.

We now prove the other direction. Suppose $\mathbb{Y} \subseteq \mathbb{R}_+^n \setminus S_n$. For any $\mathbf{q}$, $\mathbf{r} \in \mathbb{R}_+^n$, let $\mathbf{p}^* = \mathbf{p}(\mathbf{q}) = \arg\max_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot\mathbf{q} - R(\mathbf{p})$. Since $\mathbf{p}^* \in \mathbb{Y}$, $\mathbf{p}^* \cdot \mathbf{1} \geq 1$. Then $C(\mathbf{q}+\mathbf{r}) - C(\mathbf{q}) = (\max_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot(\mathbf{q}+\mathbf{r}) - R(\mathbf{p})) - (\mathbf{p}^*\cdot\mathbf{q} - R(\mathbf{p}^*)) \geq (\mathbf{p}^*\cdot(\mathbf{q}+\mathbf{r}) - R(\mathbf{p}^*)) - (\mathbf{p}^*\cdot\mathbf{q} - R(\mathbf{p}^*)) = \mathbf{p}^*\cdot\mathbf{r} \geq \min_{i\in\{1,\cdots,n\}} r_i$. □

3.3. Liquidity Adaptation

We have argued that adaptive liquidity is a valuable market property, but have not formally defined what this means. There are several reasonable ways that one could define it. Othman and Sandholm [2011] define it only informally, stating that market makers with adaptive liquidity should have "more muted price responses" at higher levels of activity. Othman and Sandholm [2012] define a similar property of "unlimited market depth" as the property that the second derivative of C with respect to i goes to 0 for all i as $||\mathbf{q}||$ grows. We avoid this definition as it requires C to be twice differentiable, which is an unnecessary restriction. Instead of looking at this "instantaneous" liquidity, we define adaptive liquidity in terms of the price change for a specific bundle $\mathbf{r}$. Our definition requires only that C be differentiable, which is already required for Existence of Instantaneous Prices.

PROPERTY 6 (LIQUIDITY ADAPTATION). *For any bundle* $\mathbf{r} \in \mathbb{R}_+^n$ *and any* $\epsilon > 0$*, there is some* $M > 0$ *such that* $\|\mathbf{q}\|_2 > M$ *implies* $\mathbf{p}(\mathbf{q}+\mathbf{r}) - \mathbf{p}(\mathbf{q}) < \epsilon$.

It turns out that there is a simple geometric characterization of duality-based buy-only markets that satisfy this property. Perhaps surprisingly, the strict non-linearity of $\partial^+\mathbb{Y}$, which is imposed in Othman and Sandholm [2011] to enforce that their homogeneous risk measures are differentiable, is both a sufficient and necessary condition.

THEOREM 3.8. *Consider a duality-based buy-only market with bounded R that satisfies Existence of Instantaneous Prices. The market is liquidity adaptive if and only if* $\partial^+\mathbb{Y}$ *is strictly non-linear.*

The technical condition for Liquidity Adaptation can be written equivalently as $\lim_{\lambda\to\infty} \mathbf{p}(\lambda\mathbf{q}+\mathbf{r}) - \mathbf{p}(\lambda\mathbf{q}) = 0$ *uniformly* for all $\mathbf{q}$ such that $\|\mathbf{q}\|_2 = 1$. We work with this alternative formulation in the proof.

We break the proof into several pieces. We first prove the sufficiency of the strictly non-linear boundary of $\mathbb{Y}$ in two steps. Lemma 3.9 shows that $\mathbf{p}(\lambda\mathbf{q})$ converges uniformly to a particular value $\hat{\mathbf{p}}(\mathbf{q})$ as λ gets big, i.e., that $\lim_{\lambda\to\infty} \mathbf{p}(\lambda\mathbf{q}) = \hat{\mathbf{p}}(\mathbf{q})$. Lemma 3.10 shows that a small perturbation on the market state $\mathbf{q}$ does not affect this convergence, and so purchasing a bundle $\mathbf{r}$ when $\|\mathbf{q}\|$ is large does not change prices. To show necessity, we first assume the existence of a line segment on the positive boundary. Then we can show that if we take $\mathbf{q}$ pointing to the direction of the norm of that line segment, a small perturbation of $\mathbf{q}$ (i.e., the purchase of a small bundle $\mathbf{r}$) can actually cause the price to move to the end of the line segment, even if $\|\mathbf{q}\|$ is large.

LEMMA 3.9. *For any duality-based buy-only market with bounded R that satisfies Existence of Instantaneous Prices, if* $\partial^+\mathbb{Y}$ *is strictly nonlinear, then* $\lim_{\lambda\to\infty} \mathbf{p}(\lambda\mathbf{q}) = \hat{\mathbf{p}}(\mathbf{q})$ *uniformly for all* $\mathbf{q}$ *s.t.* $\mathbf{q} \in \mathbb{R}_+^n$ *and* $\|\mathbf{q}\|_2 = 1$ *where* $\hat{\mathbf{p}}(\mathbf{q}) = \arg\max_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot\mathbf{q}$.

PROOF. Fix any $\mathbf{q} \in \mathbb{R}_+^n$. To simply notation, let $\hat{\mathbf{p}} = \hat{\mathbf{p}}(\mathbf{q})$. For any $\epsilon > 0$, let $B(\hat{\mathbf{p}}, \epsilon)$ be the open ball centered at $\hat{\mathbf{p}}$ with radius ϵ. By the same argument used in the proof of Proposition 3.4, $\hat{\mathbf{p}}$ must be on the positive boundary of $\mathbf{y}$. This is shown in Figure 1. By definition of $\hat{\mathbf{p}}$, for any $\mathbf{p}' \in \mathbb{Y} \setminus B(\hat{\mathbf{p}}, \epsilon)$, $\mathbf{p}'\cdot\mathbf{q} \le \hat{\mathbf{p}}\cdot\mathbf{q}$. In fact, this equation must hold with strict inequality by a similar argument to the one that shows the maximizer is unique in the proof of Proposition 3.4. Since $\mathbb{Y} \setminus B(\hat{\mathbf{p}}, \epsilon)$ is a compact set and a continuous function on a compact set attains its supremum, $\max_{\mathbf{p}'\in\mathbb{Y}\setminus B(\hat{\mathbf{p}},\epsilon)}\{\mathbf{p}'\cdot\mathbf{q}\} < \hat{\mathbf{p}}\cdot\mathbf{q}$. Therefore, we can fix some $\delta > 0$ such that

$$\max_{\mathbf{p}\in\mathbb{Y}\setminus B(\hat{\mathbf{p}},\epsilon)} \mathbf{p}\cdot\mathbf{q} < \hat{\mathbf{p}}\cdot\mathbf{q} - \delta \tag{3}$$

Let $M = \sup_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p}) - \inf_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p})$ and pick a $\lambda > M/\delta$. We will prove by contradiction that $\mathbf{p}(\lambda\mathbf{q}) \in B(\hat{\mathbf{p}}, \epsilon)$. We know $\mathbf{p}(\lambda\mathbf{q}) \in \mathbb{Y}$. Suppose $\mathbf{p}(\lambda\mathbf{q}) \in \mathbb{Y} \setminus B(\hat{\mathbf{p}}, \epsilon)$. Then $\lambda\mathbf{q}\cdot\mathbf{p}(\lambda\mathbf{q}) - R(\mathbf{p}(\lambda\mathbf{q})) < \lambda(\hat{\mathbf{p}}\cdot\mathbf{q} - \delta) - (R(\hat{\mathbf{p}}) - M) < \lambda\hat{\mathbf{p}}\cdot\mathbf{q} - M - R(\hat{\mathbf{p}}) + M = \lambda\mathbf{q}\cdot\hat{\mathbf{p}} - R(\hat{\mathbf{p}})$, which contradicts the fact that $\mathbf{p}(\lambda\mathbf{q})$ is the maximizer. Therefore $\mathbf{p}(\lambda\mathbf{q}) \in B(\hat{\mathbf{p}}, \epsilon)$.

We have shown that for any ϵ, if δ is set as above, then for any $\lambda > M/\delta$, $\|\mathbf{p}(\lambda\mathbf{q}) - \hat{\mathbf{p}}\|_2 < \epsilon$. This implies that $\lim_{\lambda\to\infty} \mathbf{p}(\lambda\mathbf{q}) = \hat{\mathbf{p}}$.

To prove that the limit process is uniform for all $\mathbf{q} \in \mathbb{R}_+^n$ such that $\|\mathbf{q}\|_2 = 1$, it is enough to find a value δ such that Equation 3 holds for all such $\mathbf{q}$. By construction, the maximum valid value of δ is a continuous function of $\hat{\mathbf{p}}(\mathbf{q})$. Notice also that $\hat{\mathbf{p}}(\mathbf{q})$ is a continuous function of $\mathbf{q}$ by strict non-linearity of $\partial^+\mathbb{Y}$. Therefore, the maximum valid value of δ is continuous in $\mathbf{q}$ and must attain a minimum value on the closed set $\{\|\mathbf{q}\|_2 = 1,\ \mathbf{q} \in \mathbb{R}_+^n\}$. This value of δ makes Equation 3 hold for all such $\mathbf{q}$. □

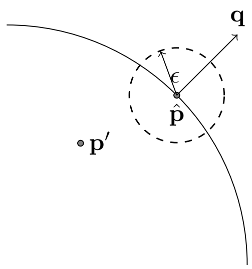

Fig. 1. Illustration of $\mathbb{Y}$ and $B(\hat{\mathbf{p}}, \epsilon)$ from the proof of Lemma 3.9.

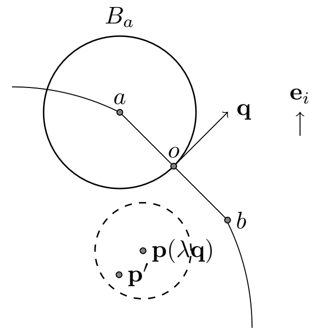

Fig. 2. Illustration of the proof of Lemma 3.11, with $\mathbf{p}(\lambda\mathbf{q})$ outside B_a.

The next lemma shows that a slight perturbation on market state $\mathbf{q}$ will not change the limit process.

Lemma 3.10. *Under the same condition as in Lemma 3.9, for any* $\mathbf{r} \in \mathbb{R}^n_+$, $\lim_{\lambda\to\infty} \mathbf{p}(\lambda\mathbf{q} + \mathbf{r}) = \hat{\mathbf{p}}(\mathbf{q})$ *uniformly for all* $\mathbf{q}$ *s.t.* $\mathbf{q} \in \mathbb{R}^n_+$ *and* $\|\mathbf{q}\| = 1$.

Proof. Fix a $\mathbf{q}$ and let $\hat{\mathbf{p}}$ denote $\hat{\mathbf{p}}(\mathbf{q})$. Let δ, M, and $B(\hat{\mathbf{p}}, \epsilon)$ be defined as in the proof of Lemma 3.9. Let $P = \sup_{\mathbf{p}\in\mathbb{Y},i} \mathbf{p} \cdot \mathbf{r}$ (which must be finite since $\mathbb{Y}$ is bounded) and pick any $\lambda > (M + P)/\delta$.

Suppose that $\mathbf{p}(\lambda\mathbf{q} + \mathbf{r}) \in \mathbb{Y} \setminus B(\hat{\mathbf{p}}, \epsilon)$. Then

$$\begin{aligned}(\lambda\mathbf{q} + \mathbf{r}) \cdot \mathbf{p}(\lambda\mathbf{q} + \mathbf{r}) - R(\mathbf{p}(\lambda\mathbf{q} + \mathbf{r})) &= \lambda\mathbf{q} \cdot \mathbf{p}(\lambda\mathbf{q} + \mathbf{r}) + \mathbf{r} \cdot \mathbf{p}(\lambda\mathbf{q} + \mathbf{r}) - R(\mathbf{p}(\lambda\mathbf{q} + \mathbf{r})) \\ &< \lambda(\mathbf{q} \cdot \hat{\mathbf{p}} - \delta) + P - (R(\hat{\mathbf{p}}) - M) < \lambda\mathbf{q} \cdot \hat{\mathbf{p}} - R(\hat{\mathbf{p}}) \\ &\le \lambda\mathbf{q} \cdot \hat{\mathbf{p}} + \mathbf{r} \cdot \hat{\mathbf{p}} - R(\hat{\mathbf{p}}) = (\lambda\mathbf{q} + \mathbf{r}) \cdot \hat{\mathbf{p}} - R(\hat{\mathbf{p}}),\end{aligned}$$

which contradicts the fact that $\mathbf{p}(\lambda\mathbf{q} + \mathbf{r})$ is the price vector at state $\lambda\mathbf{q} + \mathbf{r}$. By a similar argument to the one in the proof of Lemma 3.9, it follows that $\lim_{\lambda\to\infty} \mathbf{p}(\lambda\mathbf{q} + \mathbf{r}) = \hat{\mathbf{p}}$.

The uniformity argument is also analogous to the uniformity argument in the proof of Lemma 3.9. □

Combining this with Lemma 3.9 completes the argument that the non-linear positive boundary of $\mathbb{Y}$ is sufficient for liquidity adaptivity. Now we show it is necessary as well.

Lemma 3.11. *For any duality-based buy-only market maker with bounded* R *that satisfies Existence of Instantaneous Prices, if* $\partial^+\mathbb{Y}$ *contains a line segment, then the market does not satisfy Liquidity Adaptation.*

Proof. Suppose that $\partial^+\mathbb{Y}$ contains a line segment $\overline{\mathbf{ab}}$ between points $\mathbf{a}$ and $\mathbf{b}$ of length 2ℓ for some $\ell > 0$. Let $\mathbf{q}$ be the norm of $\mathbb{Y}$ at $\overline{\mathbf{ab}}$. Then $\mathbf{a}$, $\mathbf{b}$, and any point between them on the line segment all maximize the dot product with $\mathbf{q}$ over $\mathbb{Y}$.

Let B_a be the open ball centered at $\mathbf{a}$ with radius ℓ, as in Figure 2. Consider the function $f(\mathbf{p}) = \max_i(a_i - p_i)$. This function is continuous, and $f(\mathbf{p}) > 0$ for all $\mathbf{p} \in \mathbb{Y} \setminus B_a$ since $\mathbf{a}$ can not be dominated by any point in $\mathbb{Y}$ except itself. $\mathbb{Y} \setminus B_a$ is a compact set since $\mathbb{Y}$ is compact and B_a is an open ball. Since any continuous function achieves its supremum on a compact set, $\min_{\mathbf{p}\in\mathbb{Y}\setminus B_a} (\max_i(a_i - p_i)) = \delta_a > 0$. Similarly, if we define B_b be the open ball centered at $\mathbf{b}$ with radius ℓ, then $\min_{\mathbf{p}\in\mathbb{Y}\setminus B_b} (\max_i(b_i - p_i)) = \delta_b > 0$. Let $\delta = \min\{\delta_a, \delta_b\}$.

For any $\lambda > 0$, we know that the price vector $\mathbf{p}(\lambda\mathbf{q})$ is in $\mathbb{Y}$. If $\mathbf{p}(\lambda\mathbf{q})$ is not in the open ball B_a, then $\max_i(a_i - p_i(\lambda\mathbf{q})) \ge \delta$. Thus there is some i such that $\mathbf{a} \cdot \mathbf{e}_i \ge \mathbf{p}(\lambda\mathbf{q}) \cdot \mathbf{e}_i + \delta$. Let $\mathbf{p}'$ be any point of $\mathbb{Y}$ within $\delta/2$ distance to $\mathbf{p}(\lambda\mathbf{q})$, then $\mathbf{p}' \cdot \mathbf{e}_i \le \mathbf{p}(\lambda\mathbf{q}) \cdot \mathbf{e}_i + \delta/2 \le \mathbf{a} \cdot \mathbf{e}_i - \delta/2$.

Notice also that $\mathbf{p}' \cdot \mathbf{q} \leq \mathbf{a} \cdot \mathbf{q}$. Letting $M = \sup_{\mathbf{p} \in \mathbb{Y}} R(\mathbf{p}) - \inf_{\mathbf{p} \in \mathbb{Y}} R(\mathbf{p})$,

$$\mathbf{p}' \cdot \left(\lambda \mathbf{q} + \frac{3M}{\delta}\mathbf{e}_i\right) - R(\mathbf{p}') \leq \mathbf{a} \cdot \lambda \mathbf{q} + \frac{3M}{\delta}\left(\mathbf{a} \cdot \mathbf{e}_i - \frac{\delta}{2}\right) + M - R(\mathbf{a})$$
$$= \mathbf{a} \cdot \left(\lambda \mathbf{q} + \frac{3M}{\delta}\mathbf{e}_i\right) - M/2 - R(\mathbf{a}) < \mathbf{a} \cdot \left(\lambda \mathbf{q} + \frac{3M}{\delta}\mathbf{e}_i\right) - R(\mathbf{a}).$$

This means $\mathbf{p}'$ cannot be the price at market state $\lambda \mathbf{q} + 3M\mathbf{e}_i/\delta$, thus

$$\|\mathbf{p}\left(\lambda \mathbf{q} + (3M/\delta)\mathbf{e}_i\right) - \mathbf{p}(\lambda \mathbf{q})\|_2 > \delta/2. \tag{4}$$

A similar argument shows that Equation 4 holds for some i if $\mathbf{p}(\lambda \mathbf{q})$ is not in the open ball B_b. Since B_a and B_b don't intersect, $\mathbf{p}(\lambda \mathbf{q})$ cannot be both in B_a and B_b, so for each $\lambda > 0$ Equation 4 must hold for some i. Let L_i be the set of λ s.t. Equation 4 holds for i, then $\cup_i L_i = (0, +\infty)$. At least one of L_i must be unbounded. Choose that i, let $\mathbf{r} = (3M/\delta)\mathbf{e}_i$, and let $\epsilon = \delta/2$. The above argument shows that for any M, there exists a λ such that $\|\lambda \mathbf{q}\| > M$, but $p(\lambda \mathbf{q} + \mathbf{r}) - p(\lambda \mathbf{q}) > \epsilon$, contradicting Liquidity Adaptation. □

3.4. The Characterization Theorem

By combining the properties from the previous two sections, we can define a broad class of market makers that satisfy the six axioms we've defined.

THEOREM 3.12. *Any buy-only market maker satisfying Path Independence and Information Incorporation must be implementable via a convex cost function C. If C is additionally closed, it can be written in the form $C(\mathbf{q}) = \sup_{\mathbf{p} \in \mathbb{Y}} \mathbf{p} \cdot \mathbf{q} - R(\mathbf{p})$ for some convex set $\mathbb{Y} \subseteq \mathbb{R}^n$ and some closed and convex $R : \mathbb{Y} \to \mathbb{R}$.*

The market additionally satisfies the Existence of Instantaneous Prices if either R is strictly convex or R is affine with coefficient $\mathbf{q}_0 \in \mathbb{R}^n_+ \setminus \{\mathbf{0}\}$ and $\partial^+\mathbb{Y}$ is strictly nonlinear.

If the above properties are satisfied and R is bounded, then the market maker satisfies Bound Loss, No Arbitrage, and Liquidity Adaptation if and only if (1) $\mathbb{Y} \subseteq \mathbb{R}^n_+ \setminus \{\mathbf{x} \in \mathbb{R}^n_+ \mid \sum_{i=1}^n x_i < 1\}$; (2) $\Delta_n \subseteq \mathrm{extend}(\mathbb{Y})$; (3) $\partial^+\mathbb{Y}$ is strictly nonlinear.

This characterization is not a full "if and only if" characterization. As discussed in Section 3.2, the necessary condition for Existence of Instantaneous Prices is awkward, while the two sufficient conditions we provided are more natural and do not rule out any natural cost functions. Additionally we require explicitly that R is bounded; it might be possible to construct a market satisfying the properties without bounded R, but it would be much trickier to argue for properties like bounded loss, for which we would need to reason about exactly how $R(\mathbf{p})$ goes to infinity when $\mathbf{p}$ approaches $\partial^+\mathbb{Y}$.

In the next sections, we will explore the role of $\mathbb{Y}$ in more detail. We will see that the extent to which $\mathbb{Y}$ extends beyond the probability simplex controls the trade-off between information loss and profitability since it affects the sum of prices. Additionally, the curvature of $\mathbb{Y}$ at any point on its boundary is related to market liquidity when the price vector is near that point; this is why liquidity adaptivity requires non-linearity of the boundary.

4. INFORMATION LOSS

In the last section, we derived a class of market makers that satisfy a selection of desirable properties, summarized in Theorem 3.12. Different market makers in this class yield different properties. In this section, we consider the market's ability to aggregate information from traders and see how this is affected by the choice of market.

4.1. Sacrificing Expressiveness

The property of Abernethy et al. [2013] that is not obtained by duality-based buy-only markets is *expressiveness*. Expressiveness requires that, given sufficient wealth, a trader should be able to "push" market prices to reflect his beliefs. Specifically, if a trader's beliefs are represented by a probability distribution $\mathbf{b}$ over the outcome space, he should be able to buy or sell securities up to the point at which the price of each security is the expected payout of the security with respect to $\mathbf{b}$.[4] Of course in buy-only markets, traders can express beliefs only through purchases, not sales.

PROPERTY 7 (EXPRESSIVENESS). $\forall \mathbf{q} \in \mathbb{R}^n_+$, $\forall \mathbf{b} \in \Delta_n$, $\exists \mathbf{r} \in \mathbb{R}^n_+$ *s.t.* $\mathbf{p}(\mathbf{q}+\mathbf{r}) = \mathbf{b}$.

This property is only well-defined if the market satisfies Existence of Instantaneous Prices. Unfortunately, Expressiveness is in direct conflict with Liquidity Adaptation.

THEOREM 4.1. *A duality-based buy-only market with bounded R that satisfies Existence of Instantaneous Prices cannot satisfy Liquidity Adaptation and Expressiveness.*

PROOF. Consider a duality-based buy-only market with bounded R that satisfies Existence of Instantaneous Prices and Expressiveness. For Expressiveness to hold, it must be the case that $\Delta_n \subseteq \mathbb{Y}$. Pick any $\mathbf{x}$ in the relative interior of Δ_n, i.e., any distribution $\mathbf{x}$ with no zero component. Clearly $\mathbf{x} \in \mathbb{Y}$.

We first show that we must have $\mathbf{x} \in \partial\mathbb{Y}$. Suppose this were not the case. Then there must exist a closed ball $B(\mathbf{x}, \epsilon)$ in the interior of $\mathbb{Y}$ centered at $\mathbf{x}$ with radius $\epsilon > 0$. Recall that $\mathbb{Y}$ is the effective domain of R and recall that a convex function must be finite on any closed set contained in the relative interior of its effective domain, so R must be bounded on $B(\mathbf{x}, \epsilon)$. Let $M = \sup_{\mathbf{p} \in B(\mathbf{x},\epsilon)} |R(\mathbf{p})|$, and choose any $\mathbf{q} \in \mathbb{R}^n_+$ such that $\|\mathbf{q}\|_2 > 2M/\epsilon$. Then for any $\mathbf{r} \in \mathbb{R}^n_+$, $\mathbf{x}$ cannot maximize $\mathbf{p} \cdot (\mathbf{q}+\mathbf{r}) - R(\mathbf{p})$ on $B(\mathbf{x}, \epsilon)$, since if we choose $\mathbf{x}' = \mathbf{x} + \epsilon \cdot (\mathbf{q}+\mathbf{r})/\|\mathbf{q}+\mathbf{r}\|_2$, then $\mathbf{x}' \cdot (\mathbf{q}+\mathbf{r}) - R(\mathbf{x}') = \mathbf{x} \cdot (\mathbf{q}+\mathbf{r}) + \epsilon\|\mathbf{q}+\mathbf{r}\|_2 - R(\mathbf{x}) - [R(\mathbf{x}') - R(\mathbf{x})] \geq \mathbf{x} \cdot (\mathbf{q}+\mathbf{r}) - R(\mathbf{x}) + \epsilon\|\mathbf{q}\|_2 - 2M > \mathbf{x} \cdot (\mathbf{q}+\mathbf{r}) - R(\mathbf{x})$. Therefore by Equation 2, $\mathbf{x} \neq \mathbf{p}(\mathbf{q}+\mathbf{r})$ for all $\mathbf{r} \in \mathbb{R}^n_+$, contradicting Expressiveness.

We can conclude that $\mathbf{x} \in \partial\mathbb{Y}$. Notice that this is true for all $\mathbf{x}$ in the relative interior of Δ_n, thus $\partial\mathbb{Y}$ must includes a line segment. By Theorem 3.8, this would imply that the market does not satisfy Liquidity Adaptation. □

4.2. Bounding Information Loss Through the Price Space

Without expressiveness, there may exist market prices $\mathbf{p}$ such that two myopic and risk-neutral traders with distinct beliefs $\mathbf{b}$ and $\mathbf{b}'$ both would not trade. In this case, it is not clear what conclusion to draw about the traders' collective beliefs if the market prices converge to $\mathbf{p}$. To quantify the resulting loss of information, we can define the *belief set* $B(\mathbf{p})$ as the set of underlying beliefs with which a trader is not willing to trade if the current market price vector is $\mathbf{p}$. It is easy to argue that a myopic risk-neutral trader is not willing to trade if and only if the price for each security is at least as much as his belief for the corresponding outcome. Thus we have $B(\mathbf{p}) = \{\mathbf{b} \in \Delta_n \mid \mathbf{b} \preceq \mathbf{p}\}$.

If the market were expressive, the belief set would always contain at most a single point; traders would wish to trade whenever the prices do not perfectly reflect their beliefs. As we will see, the belief set grows as the price space $\mathbb{Y}$ deviates from the probability simplex. We would like a way to design $\mathbb{Y}$ with information loss in mind.

We take the following approach. First, we define an *information loss function* $L : \mathbb{R}^n_+ \to \mathbb{R}_+$ which defines the market's information loss given the current market prices. This function typically measures the size of the belief set at the current market prices, but can be defined in different ways depending on the goal of the market designer (see below). Based on the designer's needs, we also choose a *tolerable loss threshold* t.

Together, the information loss function and tolerable loss threshold give rise to the sub-level set $\bar{\mathbb{Y}} = \{\mathbf{p} \mid L(\mathbf{p}) \leq t\}$, the set of all price vectors that result in information loss less

[4] For technical reasons, Abernethy et al. [2013] only require that the market prices be within ϵ of the expected payouts, for arbitrarily small ϵ. This weaker property would still be in conflict with liquidity-adaptivity.

than t. If we choose $\mathbb{Y}$ to be a convex subset of $\bar{\mathbb{Y}}$, then we are guaranteed that the information loss is no more than t under the measure L, regardless of the final market state. In general we cannot set $\mathbb{Y} = \bar{\mathbb{Y}}$ because $\bar{\mathbb{Y}}$ could fail to be convex or have other bad properties. However, as we will see in the next section when we discuss profitability, it can be valuable to choose $\mathbb{Y}$ to be as large as possible while still convex.

In practice, since C is evaluated through an optimization over $\mathbb{Y}$, it is beneficial computationally if $\bar{\mathbb{Y}}$ can be represented with (a small number of) convex constraints, which is the case if the loss is convex. Keeping this in mind, we turn to some examples.

4.2.1. 1-Norm Information Loss. Suppose that an investor would like to run a market to predict whether a startup in which he might invest will succeed (in which case he will make \$$a$) or fail (in which case he will lose \$$b$). If the probability of success is p, the expected value of the investment would be $pa - (1-p)b$. If the market reports probability p', the estimated expectation would be off by $(p - p')(a + b)$, which is proportional to the 1-norm distance between the market distribution and the true distribution.

In this case, the market designer might choose the loss function $L(\mathbf{p}) = \sum_{i=1}^{n} p_i - 1$, an upper bound of the 1-norm diameter of $B(\mathbf{p})$. For a threshold t, $\bar{\mathbb{Y}} = \{\mathbf{p} \mid \sum_{i=1}^{n} p_i \leq 1 + t\}$, the set of price vectors such that the sum of prices over all securities is at most $1 + t$. Note that $\bar{\mathbb{Y}}$ is not strictly convex and extends beyond the unit hypercube for positive t, which is perhaps unnatural since traders should never purchase a security for more than \$1. One way to fix that is to take $\mathbb{Y}$ as the maximal ℓ_α-ball (not necessarily centered at the origin) tangent with $\{\mathbf{p} \mid \sum p_i = 1 + t\}$ whose surface goes through all standard basis vectors, where $\alpha > 1$ is chosen to ensure that the ball is fully contained in the unit hypercube. The constraint in the optimization problem is an ℓ_α-norm constraint, which is tractable.

4.2.2. Entropy Weighted 1-Norm Information Loss. Suppose instead that this investor is extremely risk averse. He may want to make the investment if the chance of success is near certain, but even a small chance of failure is too much. This investor would care a lot about the difference between a 1% chance of failure and a 5% chance, but wouldn't care much about the difference between a 20% chance and an 80% chance since he won't make the investment either way. In this scenario, the 1-norm loss does not make sense. Instead, we might consider an *entropy-weighted* 1-norm loss, defined as $L(\mathbf{p}) = (\sum_{i=1}^{n} p_i - 1)\log(n)/H(\mathbf{p}_\pi)$, where H is the entropy function and $\mathbf{p}_\pi$ is the projection of $\mathbf{p}$ onto the hyperplane $\{\mathbf{p} \mid \sum_{i=1}^{n} p_i = 1\}$, used in place of $\mathbf{p}$ so that entropy is well-defined. The factor $\log(n)$ simply normalizes the weights into the range $[1, +\infty)$. This loss function captures the intuition that when the entropy is small, the market prediction is in the valuable range, and we would like the price vector to be close to the probability simplex to make more accurate prediction; when entropy is high, the prediction is not valuable anyway so we can afford to let the price drift further from the simplex (which will result in higher profit as we'll see in Section 5).

For fixed $t \geq 0$, the constraint $(\sum_{i=1}^{n} p_i - 1)\log(n)/H(\mathbf{p}_\pi) \leq t$ can be written as $-t(\log(n))^{-1}H(\mathbf{p}_\pi) + \sum_{i=1}^{n} p_i - 1 \leq 0$, which is a convex constraint. However, it is possible, as we can see from Figure 3, that $\bar{\mathbb{Y}}$ can go out of the unit hypercube for large t. In that case, we could let $\mathbb{Y} = \bar{\mathbb{Y}} \cap B_\alpha(0, 1)$, where $B_\alpha(c, r)$ is the ℓ_α-ball centered at c with radius r. Notice that $B_\alpha(0, 1)$ is an approximation of unit hypercube for large α. For the optimization problem, we can simply add the ℓ_α-norm constraint $\|\mathbf{p}\|_\alpha \leq 1$. In practice, t would likely be small enough that this would not be necessary.

The information loss we have described is different from the notion of information loss discussed by Abernethy et al. [2013]. They consider loss due to the fact that in practice there is a minimum trade unit, resulting in a spread even if the *instantaneous* buy and sell prices are identical. In our framework, the instantaneous price and unit purchase price are identical for large trading volumes, so this problem goes away.

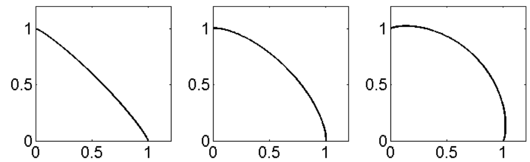

Fig. 3. Boundary of the sublevel set $\bar{\mathbb{Y}}$ defined by the entropy weighted 1-norm loss, with $t = 0.1, 0.3, 0.5$.

5. PROFITABILITY

All duality-based buy-only markets with $\mathbb{Y}$ and R satisfying the constraints in Theorem 3.12 guarantee bounded worst case loss for the market maker. However, in some cases, the market maker may be able to guarantee a profit no matter which outcome occurs. In this section, we consider the circumstances under which the market maker might be guaranteed a profit, and see how the choice of $\mathbb{Y}$ and R affects this guarantee.

5.1. When to Expect a Profit

We begin by discussing several necessary conditions for guaranteed profit. The first two conditions, sufficient volume of trade and disagreement among the traders, are properties of the traders in the market and their beliefs rather than the market mechanism itself, and therefore cannot be controlled directly by the market designer. The third, strong convexity of the price space $\mathbb{Y}$, is under the designer's control.

5.1.1. Sufficient Trade Volume. It is not possible to guarantee that a market will be profitable if the volume of trade is low. As an extreme example, suppose there is only a single trader in the market. If the market maker is guaranteed a profit, the trader clearly has no incentive to participate in the market, and trade will not occur. In order to incentivize trade at the start of the market, the market maker must take some risk.

However, our markets could be profitable when $\|\mathbf{q}\|$ is sufficiently large. Intuitively, one can think of profit arising due to the prices of the securities summing to something more than one, in which case we can think of the market maker charging a fee for each transaction. In our markets, this happens when $\mathbf{p}$ approaches the boundary of $\mathbb{Y}$.

5.1.2. Disagreement Among the Traders. We also cannot expect a market to be profitable if traders generally agree on the correct outcome. Consider again a market predicting the winner of Best Picture in the Academy Awards. If all traders believe that *Argo* will win and buy only the *Argo* security, the market maker cannot hope to make a profit in the event that they're correct.

We formalize this idea in terms of the quantity $d(\mathbf{p}) := \min_{i\in\{1,\cdots,n\}} \|\mathbf{p} - \mathbf{e}_i\|$, where $\mathbf{p}$ is the current market price vector. In the situation described above, $\mathbf{p}$ will be close to the vector $\mathbf{e}_i$ where i represents outcome *Argo*, so $d(\mathbf{p})$ will be small. It will be larger when prices are less skewed. The profit bounds we provide depend on this quantity.

5.1.3. Strong Convexity of the Price Space. The final condition that we need to guarantee a profit is more technical and specific to our framework, but under the control of the market designer. Intuitively, profitability comes from the fact that the price space $\mathbb{Y}$ curves outwards from Δ_n, resulting in prices that sum to more than one. Theorem 3.12 requires that $\mathbb{Y}$ has a strictly non-linear boundary, but that is not enough to guarantee profitability; indeed, we could design a space $\mathbb{Y}$ satisfying this property that is arbitrary close to Δ_n. Instead, we require that $\mathbb{Y}$ be *strongly convex* [Vial 1983]. Strong convexity can be defined in different ways. Informally, a set $S \subseteq \mathbb{R}^n$ is strongly convex with radius $r > 0$ if for any $\mathbf{x} \in \partial S$, $S \subseteq B_r(\mathbf{x})$, where H is the supporting hyperplane of S at $\mathbf{x}$ and $B_r(\mathbf{x})$ is the n-dimensional closed Euclidean ball with radius r tangent to $\mathcal{H}$ at $\mathbf{x}$ on the same side of $\mathcal{H}$ as S. This

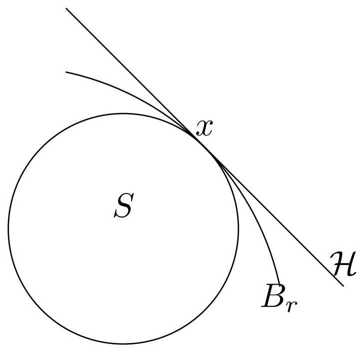

Fig. 4. Illustration of strong convexity. S is strongly convex with radius r.

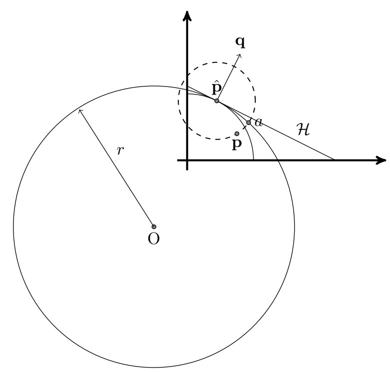

Fig. 5. Cross section of the tangent ball B_R at $\bar{\mathbf{p}}$ along its great circle.

definition is illustrated in Figure 4. If S is strongly convex with radius r, we can say it has *curvature* at least $1/r$.

The markets considered in the previous section have strongly convex price spaces. If a market is defined in terms of an ℓ_2-ball, then r is the radius of that ball. Markets defined in terms of ℓ_α balls for other α are strongly convex too, though r is not as simple. Our profit bounds hold even when r (or equivalently, the curvature of $\mathbb{Y}$) is not known exactly.

5.2. Bounding the Profit

The following theorem shows that when the three conditions described above are satisfied, our framework allows the market maker to make a guaranteed profit, regardless of the outcome. The left-hand side of Equation 5 can be interpreted as the "profit rate" of the market at its final state $\mathbf{q}$, the ratio between the minimum profit (independent of the outcome) and the amount of money collected from traders. The bound shows a clear tradeoff between the quality of collected information (as measured by the disagreement $d(\mathbf{p})$), loss of information (measured by r, since a smaller r leads to prices further from the simplex and therefore higher information loss), and profitability.

THEOREM 5.1. *For any duality-based buy-only market with a price space $\mathbb{Y}$ that is strongly convex with radius r and contains all standard basis vectors, if the final market price is $\mathbf{p}$, and the final market state $\mathbf{q}$ satisfies $\|\mathbf{q}\| \geq 16rM/d(\mathbf{p})^2$ where $M = \sup_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p}) - \inf_{\mathbf{p}\in\mathbb{Y}} R(\mathbf{p})$, we have*

$$\frac{C(\mathbf{q}) - C(\mathbf{0}) - \max_{i\in\{1,\cdots,n\}} q_i}{C(\mathbf{q}) - C(\mathbf{0})} \geq \frac{d(\mathbf{p})^2}{d(\mathbf{p})^2 + 16r}. \tag{5}$$

PROOF. Let $\hat{\mathbf{p}} = \arg\max_{\mathbf{p}\in\mathbb{Y}} \mathbf{p}\cdot\mathbf{q}$. We must have that $\hat{\mathbf{p}} \in \partial^+\mathbb{Y}$. We first show that $\|\mathbf{p} - \hat{\mathbf{p}}\| \leq d(\mathbf{p})/2$. For contradiction, suppose this were not the case.

Since $\mathbb{Y}$ is strongly convex with radius r, we know that there exists a ball $B_r(\hat{\mathbf{p}})$ of radius r tangent to the supporting hyperplane H of $\mathbb{Y}$ at $\hat{\mathbf{p}}$ such that $\mathbb{Y} \subseteq (B_r(\hat{\mathbf{p}}))$. Figure 5 shows a cross section of $B_r(\hat{\mathbf{p}})$ along its great circle.

Draw another ball with radius $\|\mathbf{p} - \hat{\mathbf{p}}\|$ centered at $\hat{\mathbf{p}}$, and let $\mathbf{a}$ be a point where it intersects $B_r(\hat{\mathbf{p}})$ as in the figure. Let $d(\mathbf{a}, H)$ and $d(\mathbf{p}, H)$ be the distances from $\mathbf{a}$ and $\mathbf{p}$ to H respectively. By similarity of triangles, we can get $d(\mathbf{a}, H) = \|\hat{\mathbf{p}} - \mathbf{a}\|^2/(2r)$. To see this, let us look at Figure 6, which is a locally magnified version of Figure 5. In this figure, we draw the diameter of the circle from $\hat{\mathbf{p}}$ to its opposite point t, we draw $\overline{\mathbf{as}}$ orthogonal to H, and we connect $\overline{\hat{\mathbf{p}}\mathbf{t}}$, $\overline{\hat{\mathbf{p}}\mathbf{a}}$ and $\overline{\mathbf{at}}$. Notice that $\angle \mathbf{ta}\hat{\mathbf{p}} = \angle \mathbf{as}\hat{\mathbf{p}} = \angle \mathbf{s}\hat{\mathbf{p}}\mathbf{t} = 90°$, thus $\angle \mathbf{a}\hat{\mathbf{p}}\mathbf{t} = 90° - \angle \mathbf{s}\hat{\mathbf{p}}\mathbf{a}$ and so $\triangle \mathbf{ta}\hat{\mathbf{p}} \sim \triangle \hat{\mathbf{p}}\mathbf{sa}$. We then have $\|\mathbf{a} - \mathbf{s}\|/\|\hat{\mathbf{p}} - \mathbf{a}\| = \|\hat{\mathbf{p}} - \mathbf{a}\|/\|\hat{p} - \mathbf{t}\|$. Notice that $\|\mathbf{a} - \mathbf{s}\| = d(\mathbf{a}, H)$, $\|\hat{\mathbf{p}} - \mathbf{t}\| = 2r$, thus $d(\mathbf{a}, H) = \|\hat{\mathbf{p}} - \mathbf{a}\|^2/(2r)$.

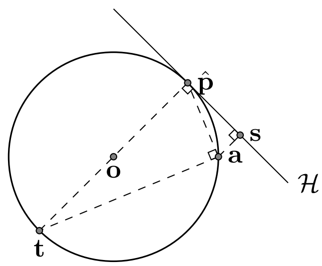

Fig. 6. Local magnification of Figure 5.

Fig. 7. The trade-off between profit rate and information loss for Euclidean ball market makers.

Since $\|\hat{\mathbf{p}} - \mathbf{a}\| = \|\hat{\mathbf{p}} - \mathbf{p}\| > d(\mathbf{p})/2$, we have $d(\mathbf{a}, H) > d(\mathbf{p})^2/(8r)$. By basic geometry, $d(\mathbf{p}, H) \geq d(\mathbf{a}, H)$, so $d(\mathbf{p}, H) > d(\mathbf{p})^2/(8r)$.

Let $\bar{\mathbf{q}} = \mathbf{q}/\|\mathbf{q}\|$. Since $\bar{\mathbf{q}}$ is the norm of H and $\hat{\mathbf{p}}$ is on H, $(\hat{\mathbf{p}} - \mathbf{p}) \cdot \bar{\mathbf{q}} = d(\mathbf{p}, H)$. We have $(\hat{\mathbf{p}} \cdot \mathbf{q} - R(\hat{\mathbf{p}})) - (\mathbf{p} \cdot \mathbf{q} - R(\mathbf{p})) = \|\mathbf{q}\|(\hat{\mathbf{p}} - \mathbf{p}) \cdot \bar{\mathbf{q}} - (R(\hat{\mathbf{p}}) - R(\mathbf{p})) > 8rM/d(\mathbf{p})^2 \cdot d(\mathbf{p})^2/(8r) - M = 0$, which contradicts the fact that $\mathbf{p} = \arg\max_{\mathbf{p} \in \mathbb{Y}} \mathbf{p} \cdot \mathbf{q} - R(\mathbf{p})$. Therefore $\|\mathbf{p} - \hat{\mathbf{p}}\| \leq d(\mathbf{p})/2$.

Let $k = \arg\max_{i \in \{1, \cdots, n\}} q_i$. Since $\|\mathbf{p} - \mathbf{e}_k\| \geq d(\mathbf{p})$ and $\|\mathbf{p} - \hat{\mathbf{p}}\| \leq d(\mathbf{p})/2$, it must be the case that $\|\hat{\mathbf{p}} - \mathbf{e}_k\| \geq d(\mathbf{p})/2$. Note $\mathbf{e}_i \in \mathbb{Y}$ by assumption, so e_i is just another price vector. From the same argument we used to show that $d(\mathbf{p}, H) \geq d(\mathbf{p})^2/(8r)$ above, we can get $d(\mathbf{e}_k, H) \geq d(\mathbf{p})^2/(8r)$, so $d(\mathbf{e}_k, H) = (\hat{\mathbf{p}} - \mathbf{e}_k) \cdot \bar{\mathbf{q}} = \hat{\mathbf{p}} \cdot \bar{\mathbf{q}} - \bar{q}_k \geq d(\mathbf{p})^2/(8r)$. Finally,

$$
\begin{aligned}
\frac{C(\mathbf{q}) - C(\mathbf{0}) - q_k}{C(\mathbf{q}) - C(\mathbf{0})} &= 1 - \frac{q_k}{\mathbf{p} \cdot \mathbf{q} - R(\mathbf{p}) - C(\mathbf{0})} \geq 1 - \frac{q_k}{\hat{\mathbf{p}} \cdot \mathbf{q} - R(\hat{\mathbf{p}}) + \sup_{\mathbf{y} \in \mathbb{Y}}(-R(\mathbf{y}))} \\
&\geq 1 - \frac{q_k}{\hat{\mathbf{p}} \cdot \mathbf{q} - M} = 1 - \frac{\bar{q}_k}{(\hat{\mathbf{p}} \cdot \bar{\mathbf{q}} - \bar{q}_k) - \frac{M}{\|\mathbf{q}\|} + \bar{q}_k} \geq 1 - \frac{\bar{q}_k}{\frac{d(\mathbf{p})^2}{8r} - \frac{M}{\|\mathbf{q}\|} + \bar{q}_k} \\
&\geq 1 - \frac{\bar{q}_k}{\frac{d(\mathbf{p})^2}{8r} - \frac{d(\mathbf{p})^2}{16r} + \bar{q}_k} = \frac{\frac{d(\mathbf{p})^2}{16r}}{\frac{d(\mathbf{p})^2}{16r} + \bar{q}_k} = \frac{d(\mathbf{p})^2}{d(\mathbf{p})^2 + 16r}.
\end{aligned}
$$

□

It can be more convenient to measure actual profit. If $\mathbb{Y}$ contains all standard basis vectors, then $C(\mathbf{q}) - C(\mathbf{0}) \geq \max_{\mathbf{p} \in \mathbb{Y}} \mathbf{p} \cdot \mathbf{q} - M \geq \|\mathbf{q}\|_\infty - M$. The previous theorem then implies the following. Notice the dependence on $\|\mathbf{q}\|_\infty$, a measure of trade volume.

Corollary 5.2. *Under the same conditions as Theorem 5.1, if $\|\mathbf{q}\|_\infty > M$, then the profit for the market maker is as least $(\|\mathbf{q}\|_\infty - M)d(\mathbf{p})^2/(d(\mathbf{p})^2 + 16r)$.*

5.3. Trading Off Information and Profit

We have seen that when the boundary of the price space moves further from the simplex, the sum of prices grows and profit increases. However, this also results in larger belief sets, leading to increased information loss. Thus there is a clear tradeoff between profit and information loss. There is no one "optimal" market; a market designer who cares about both quantities must determine how to balance them.

As an example of how this can be done, suppose that the designer cares about 1-norm information loss and profit rate, and would like to implement a market from the class of Euclidean ball market makers, where $\mathbb{Y}$ is part of a Euclidean ball with radius $r \geq 1$ whose surface crosses all standard basis vectors. Let $n = 2$ and suppose the designer speculates that the difference in beliefs ($d(\mathbf{p})$) will be at least 0.1 (though of course he cannot know for sure). In this case,the maximum information loss can be calculated as $\max_{\mathbf{p}\in\mathbb{Y}} \sum p_i - 1 = \sqrt{2}(r - \sqrt{r^2 - 0.5})$. Assuming $\|\mathbf{q}\|$ is large enough to satisfy the requirements of Theorem 5.1, the profit rate bound is $1/(1 + 1600r)$. Figure 7 illustrates the resulting tradeoff between profit rate and information loss. We can see that in the limit (bottom-left corner of the curve), information loss and profit both go to zero.

6. CONCLUSION

We derived an axiomatic characterization of a parameterized class of market makers with adaptive liquidity. We showed that the key to adaptive liquidity is a price space that extends beyond the probability simplex, and that properties of this price space (such as its curvature) can be used to quantify trade-offs between the market's ability to incorporate information about traders' beliefs and the potential for profitability.

This research suggests several interesting avenues for future work. First, our characterization applies only to *complete* markets, whereas the characterization of Abernethy et al. [2013] also applied to *complex* markets in which security payoffs can be arbitrary. While we believe it is possible to extend the ideas in this paper to that setting, our geometric proofs would not carry over easily. Second, all of our properties were defined in terms of cost functions. One could derive a characterization of liquidity-adaptive markets that do not satisfy path independence and therefore cannot be expressed in the cost function framework.

ACKNOWLEDGMENTS

We are grateful to Jake Abernethy and Yiling Chen for early discussions on this topic, to Jake Abernethy and Raf Frongillo for enlightening conversations about alternative notions of information loss, and to Lieven Vandenberghe for useful feedback on an earlier draft of this work.

REFERENCES

Abernethy, J., Chen, Y., and Wortman Vaughan, J. 2011. An optimization-based framework for automated market-making. In *ACM EC*.

Abernethy, J., Chen, Y., and Wortman Vaughan, J. 2013. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*. To appear.

Agrawal, S., Delage, E., Peters, M., Wang, Z., and Ye, Y. 2011. A unified framework for dynamic prediction market design. *Operations Research 59,* 3, 550–568.

Chen, Y. and Wortman Vaughan, J. 2010. A new understanding of prediction markets via no-regret learning. In *ACM EC*.

Hanson, R. 2003. Combinatorial information market design. *Info. Systems Frontiers 5,* 1, 105–119.

Hanson, R. 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets 1,* 1, 3–15.

Othman, A. and Sandholm, T. 2011. Liquidity-sensitive automated market makers via homogeneous risk measures. *WINE*.

Othman, A. and Sandholm, T. 2012. Profit-charging market makers with bounded loss, vanishing bid/ask spreads, and unlimited market depth. In *ACM EC*.

Othman, A., Sandholm, T., Pennock, D., and Reeves, D. 2010. A practical liquidity-sensitive automated market maker. In *ACM EC*.

Pennock, D. M. 2010. Why automated market makers? Oddhead Blog. Accessed at `http://blog.oddhead.com/2010/07/08/why-automated-market-makers/`.

Pennock, D. M. and Sami, R. 2007. Computational aspects of prediction markets. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Eds. Cambridge U. Press.

Rockafellar, R. 1997. *Convex analysis*. Vol. 28. Princeton U. Press.

Vial, J. 1983. Strong and weak convexity of sets and functions. *Mathematics of Operations Research 8,* 2, 231–259.