

A Practical Liquidity-Sensitive Automated Market Maker

ABRAHAM OTHMAN, Carnegie Mellon University
 DAVID M. PENNOCK, Microsoft Research
 DANIEL M. REEVES, Beeminder
 TUOMAS SANDHOLM, Carnegie Mellon University

Automated market makers are algorithmic agents that enable participation and information elicitation in electronic markets. They have been widely and successfully applied in artificial-money settings, like some Internet prediction markets. Automated market makers from the literature suffer from two problems that contribute to their impracticality and impair their use beyond artificial-money settings: first, they are unable to adapt to liquidity, so that trades cause prices to move the same amount in both heavily and lightly traded markets, and second, in typical circumstances, they run at a deficit. In this article, we construct a market maker that is both sensitive to liquidity and can run at a profit. Our market maker has bounded loss for any initial level of liquidity and, as the initial level of liquidity approaches zero, worst-case loss approaches zero. For any level of initial liquidity we can establish a boundary in market state space such that, if the market terminates within that boundary, the market maker books a profit regardless of the realized outcome.

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

General Terms: Economics

Additional Key Words and Phrases: Market design, automated market making, prediction markets, pricing rules, market scoring rules, liquidity

ACM Reference Format:

Othman, A., Pennock, D. M., Reeves, D. M., and Sandholm, T. 2013. A practical liquidity-sensitive automated market maker. *ACM Trans. Econ. Comp.* 1, 3, Article 14 (September 2013), 25 pages.
 DOI: <http://dx.doi.org/10.1145/2509413.2509414>

1. INTRODUCTION

Active markets like the New York Stock Exchange provide two benefits: (1) price takers can buy or sell at any time, and (2) observers can continually monitor precise values of every asset. A prediction market, or any market explicitly designed to uncover the value of an asset, relies heavily on (2) holding true. If an asset has poor price support (i.e., no open interest, or large bid-ask spread), then observers learn little or nothing about its value, disabling the very purpose of the market. For example, some popular contracts on intrade.com, one of the largest prediction markets, attract millions of

A. Othman and T. Sandholm were supported by the National Science Foundation under grants CCF-1101668, IIS-0905390, and IIS-0964579. A. Othman was supported by the Google Fellowship in Market Algorithms.

An early version of this article, under the same title, appeared in *Proceedings of the 2010 ACM Conference on Electronic Commerce*.

Authors' addresses: A. Othman and T. Sandholm, Computer Science Department, Carnegie Mellon University; email: {aothman, sandholm}@cs.cmu.edu; D. M. Pennock, Microsoft Research, New York; email: dpennock@microsoft.com; D. M. Reeves, Beeminder; email: dreeves@beeminder.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 2167-8375/2013/09-ART14 \$15.00

DOI: <http://dx.doi.org/10.1145/2509413.2509414>

dollars in trades. But thousands of other Intrade contracts suffer from low liquidity and thus reveal little in the way of predictive information.

Prediction markets can therefore benefit from an automated market maker: an algorithmic trader that always stands ready to interact with traders, providing liquidity that may be hard to support organically. For more complex environments, automated market making becomes a necessity—combinatorial prediction markets with vast numbers of outcomes to predict (e.g., a 64-team tournament with 2^{63} or 9.2 quintillion outcomes) are essentially unusable without some form of automated pricing.

Internet prediction markets are just one application of automated market making. The market makers we describe here are appropriate for use with any assets that trade off a *binary payoff* structure, in which the future can be partitioned into a finite number of states exactly one of which will be realized. For instance, companies like WeatherBill (weather insurance) and Bet365 (sports betting) are beginning to use proprietary automated market makers to offer instantaneous price quotes across thousands or millions of highly customizable assets. These kinds of binary payout structures are also becoming more prominent within traditional finance. For instance, the Chicago Board Options Exchange (CBOE) now offers binary options on the S&P and Volatility indices. While currently lightly traded relative to standard options, their integration into the largest options exchange in the U.S. augurs well for their future. Credit default swaps (CDS), which resemble insurance on bonds, have this kind of binary payout structure as well, in which the underlying bond either experiences a default event or does not. The total size of the CDS market was recently estimated at about 28 trillion dollars, making it one of the largest markets in the world [Williams 2009].

The most popular automated market maker used in Internet prediction markets is Hanson's logarithmic market scoring rule (LMSR), an automated market maker with particularly desirable properties [Hanson 2003, 2007]. The LMSR is used by a number of companies including Inkling Markets, Consensus Point, Yahoo!, Microsoft, and the large-scale non-commercial Gates Hillman Prediction Market at Carnegie Mellon [Othman and Sandholm 2010a]. (Other companies like HSX.com and Crowdcast employ their own automated market makers.) The LMSR is also the focus of academic studies about market microstructure [Ostrovsky 2009; Othman and Sandholm 2010b] and laboratory studies of market maker performance [Das 2008].

The amount of liquidity in the LMSR is a parameter set *a priori* before the market maker knows what bets traders will place. Setting the liquidity is more art than science—a constant dilemma for almost everyone who has implemented the LMSR. For instance, in the Gates Hillman Prediction Market [Othman and Sandholm 2010a], the amount of liquidity was set too low, which caused problems for traders in practice in the later stages of the market. Too little liquidity makes prices fluctuate wildly after every trade; too much makes prices barely budge even following large bets. Exacerbating the problem, the amount prices move for a fixed bet in the LMSR is a constant. The billionth-and-first dollar moves prices as much as the first. This is not the way real money markets behave; heavily traded assets like popular equities have vanishing bid/ask spreads and the ability to enter or exit large positions without significantly impacting prevailing prices, while lightly traded assets like boutique bond issues have enormous trading costs associated with them.

Liquidity is good for traders but comes at the cost of increasing the market maker's worst-case loss. In general, an LMSR operator can expect to lose money in proportion to the liquidity it provides [Hanson 2007; Pennock and Sami 2007]. The cost is rationalized as payment for traders' information. But in the real world, the vast majority of market makers run at a profit. It is no coincidence that most examples of LMSR in practice are games based on virtual currency rather than real money.

In this article, we present a variant of the LMSR that is better suited for practical use in two ways. First, our market maker automatically adjusts how easily prices change according to how much activity it sees: prices become less elastic as more dollars flow in. The market operator need not somehow try to anticipate traders' level of interest to set liquidity manually. Second, our market maker can ensure an arbitrarily small loss in the worst case and a positive profit over a wide range of final states. In the LMSR, prices of disjoint assets always sum to exactly \$1. In our market maker, prices can sum to greater than \$1. However, we prove that dropping the sums-to-unity property is a theoretical requirement for any liquidity-sensitive and path-independent market maker. Moreover, relaxing this property is precisely what allows our market maker to expect a profit, more closely resembling the market makers we see used in practice. Furthermore, we are able to obtain these properties while retaining an explicit, easy-to-calculate functional form for our market maker—one of the characteristics that makes the LMSR so popular.

Increasing market depth with increased trade may not be appropriate in every setting. Consider a market with capital-constrained traders where the true state of the world fluctuates frequently. In this setting, a constant shallow amount of market depth will allow traders to quickly reach the true state of the world. In contrast, increasing market depth with transaction volume in these settings will result in “sticky” prices that are unable to reach their correct values. However, a fluctuating true state of the world does not necessarily pose a problem for our new market maker; if the trading population is not capital constrained, prices could still be changed to reflect their putatively proper values. So, for settings in which new information does not emerge, where information is revealed gently, or where there is substantial capital “on the sidelines” waiting for trading opportunities to arise, our new market maker can obviate the need to correctly select the liquidity parameter in the LMSR.

In Section 2, we motivate the properties of automated market makers from first principles using vector calculus. We show that no market maker can satisfy three desirable properties: path independence, translation invariance, and liquidity sensitivity. With this motivation, in Section 3, we introduce our market maker, which weakens the property of translation invariance. We illustrate the features of our market maker in detail in Section 4, including obtaining tight bounds on the sum of prices.

2. PRICING RULES

A pricing rule calculates the prices that an automated market maker offers to traders. In this section, we derive from first principles the properties of pricing rules from vector calculus. This study will allow us to explore the central tension behind automated market making: that no market maker can be liquidity sensitive, path independent, and translation invariant. This axiomatic characterization is distinct from the work of Chen and Pennock [2007], who explore utility-based market makers; Agrawal et al. [2009], who use convex optimization to synthesize different strands of automated market making; Chen and Vaughan [2010], who explore the relation between no-regret learning and automated market makers; and Abernethy et al. [2011] who develop an axiomatic approach to market making focusing on combinatorial markets.

2.1. Vector Calculus for Pricing Rules

We begin by partitioning the event space into n distinct exhaustive events, exactly one of which will occur. The state of the market is kept by the quantity vector \mathbf{q} , whose i th element determines the payout owed to traders if the i th event occurs. The market maker fields bets from traders, observes the event that happens, and then settles those bets with the traders.

For instance, imagine that a market maker is taking bets on whether the Yankees or Red Sox will win in their next baseball match. A market maker with $\mathbf{q} = (1, 2)$ will pay out one dollar to traders if the Yankees win, and two dollars to traders if the Red Sox win. In automated market makers, the marginal prices of each event are a function of the obligations the market maker owes. For instance, in our example the marginal price of the Yankees winning (the first event) might be 0.4, and the marginal price of the Red Sox winning (the second event) might be 0.6. These marginal prices are the instantaneous cost of accumulating a payout on each event. As traders place bets with the market maker, the market maker's \mathbf{q} will change, which could change the marginal prices offered by the market maker. A *pricing rule* translates between quantity vectors and marginal prices.

Definition 2.1. A *pricing rule* is a differentiable function $\mathbf{p} : \mathbb{R}^n \mapsto [0, 1]^n$ that maps a vector of quantities to a vector of prices.

Pricing rules should satisfy a further property: that they have a *convex pre-image*. (However, we do not require that the pre-image of the pricing rule encompasses the entire domain \mathbb{R}^n .) Convexity is a natural property. Imagine a trader holding a portfolio \mathbf{q} . Convexity ensures that the trader can sell any fraction of that portfolio back to the market maker and still have defined prices. We now define this notion formally.

Definition 2.2. A point in \mathbb{R}^n is *valid* if it is in the pre-image of the pricing rule \mathbf{p} .

Definition 2.3. Pricing rule \mathbf{p} has a *convex pre-image* if all convex combinations of valid vectors are also valid.

Throughout the rest of the article, we will assume all pricing rules have a convex pre-image.

2.1.1. Three Desirable Properties. We can identify three desirable properties one would like a pricing rule to have: that it be *path independent*, that it be *translation invariant*, and that it be *liquidity sensitive*.

Path independence means that any way the market moves from one state to another state yields the same payment or cost to the traders in aggregate [Hanson 2003].

Definition 2.4 (Path Independence). Pricing rule \mathbf{p} is *path independent* if the value of line integral (cost) between any two quantity vectors depends only on those quantity vectors, and not on the path between them.

Path independence offers three important benefits. First, it is a sufficient condition for ensuring that there does not exist a money pump in the market: a trader cannot place a series of trades and profit without assuming some risk. Second, it provides a minimum representation of state: we only need to know the quantity vector. Finally, because a trader gets the same odds from participating all at once as in a set of small trades, traders do not need to strategize how they make trades (e.g., making a series of small purchases instead of a single large trade).

Path independence also follows from interpreting market makers as ways of assessing the riskiness of a distribution of holdings. A recent stream of research [Agrawal et al. 2009; Ben-Tal and Teboulle 2007; Othman and Sandholm 2011] has fleshed out the correspondence between cost-function automated market makers and *risk measures* from the finance literature [Artzner et al. 1999; Carmona 2009; Carr et al. 2001; Föllmer and Schied 2002]. Risk measures are frequently used as internal tools within financial institutions to determine the exposure and quality of positions. In short, risk measures take as input a vector of holdings and produce a judgment as to whether that vector is acceptable or not. Risk measures are naturally path independent, because for internal risk measurement purposes it does not matter how a portfolio was obtained—what matters is what the portfolio consists of. Put another way, it does not matter if a

contract was inherited from a legacy company, purchased at discounted price, or is in the middle of an orderly wind-down; the company holding that contract is still exposed to it identically. Because risk measures assess the quality of a vector of payouts without regard to the path taken to produce that vector, they are path independent.

Now, an important connection follows immediately from vector calculus:

LEMMA 2.5. *If a pricing rule \mathbf{p} is path independent and has a convex pre-image, then \mathbf{p} is the gradient of a scalar potential field.*

Tying this to convention in the prediction market literature, we call this scalar field a *cost function* and denote it by $C(\cdot)$. The cost function maps vectors of quantities to a single scalar value, and prices are determined by the partial derivatives with respect to each coordinate of the vector. To move from obligation vector \mathbf{x} to obligation vector \mathbf{y} , the trader pays the market maker $C(\mathbf{y}) - C(\mathbf{x})$, where negative values indicate the market maker paying the trader. Recalling our example of the Yankees-Red Sox baseball game, a trader that wishes to move the market maker's quantity vector from $(1, 2)$ to $(2, 2)$ (i.e., placing a bet that pays out one dollar if the Yankees win) would pay $C((2, 2)) - C((1, 2))$ to the market maker.

The cost function represents the (path independent) integral over instantaneous prices, so it is a measure of how much money has been paid into the system. To view this another way, imagine that a set of traders, collectively, has d dollars and the market is initially at state $C(\mathbf{q}^0)$. After all the traders invest all their money, the combined holdings of the traders can be those vectors \mathbf{q} such that

$$C(\mathbf{q}^0 + \mathbf{q}) = C(\mathbf{q}^0) + d$$

The second desired property is *translation invariance* [Agrawal et al. 2009]: that the cost of buying a guaranteed payout of x always costs x . Translation invariance has been a standard feature of market makers in the academic literature [Agrawal et al. 2009; Chen and Vaughan 2010; Hanson 2003, 2007; Pennock and Sami 2007].

Definition 2.6 (Translation Invariance). A pricing rule is *translation invariant* if prices always sum to unity. Formally:

$$\sum_i p_i(\mathbf{q}) = 1$$

for all valid \mathbf{q} .

Most markets in practical use do not preserve translation invariance, and with good reason: a translation invariant pricing rule ensures that the market maker will take a loss as long as the final market prices are more accurate than the initial market prices, a condition that is essentially tautological (if it were false, there would be little reason to run a market in the first place). The simplest way to see this is to characterize the way market makers function in standard, familiar markets. A market maker takes on a risk when setting prices: if the prices are not the actual expected final prices, the market maker has a negative expectation. Market makers counter this risk by charging different prices on both market sides (buying and selling) so that buying a guaranteed payoff of one dollar costs more than a dollar. Then, the market maker profits from traders purchasing on both sides of the market, leaving a cut (aka the "spread" or "vig") for the market maker. A translation invariant rule shrinks the size of the spread to zero, leaving the market maker exposed to the negative downside risk of offering prices without any upside.

Conversely, the translation invariance condition guarantees that no trader can arbitrage (exploit without risk) the market maker by taking on a guaranteed payout for less than the payout.

The most direct benefit of a translation invariant pricing rule is that it preserves the equality between the *price* of an event and the *probability* of that event occurring. Both prices and probabilities will be non-negative and (when exhaustively partitioned) will sum to one.

Translation invariant rules also guarantee the “law of one price”, so that if two bets offer the same payouts in all states, they will have the same price. Put another way, recall our Yankee-Red Sox baseball game example. The law of one price asserts that placing a bet of a certain amount towards the Yankees winning will be priced the same as placing a bet of the same amount on the Red Sox losing. While logically straightforward, this condition does not necessarily hold in practice in traditional continuous double auctions, as the administrators of the Iowa Electronic Markets have discussed [Berg et al. 2001; Oliven and Rietz 2004]. This condition can thus also be viewed as a necessary condition for efficient information aggregation in a market. If the law of one price is not satisfied, there are opportunities for unsophisticated traders to pay too much or get paid too little.

As the third desired property, we would like market makers to adjust the elasticity of their pricing response based on the volume of activity in the market. We call market makers that are unable to adjust in this way *liquidity insensitive*.

Definition 2.7 (Liquidity Sensitivity). Define the n -dimensional vector $\mathbf{1} \equiv (1, 1, \dots, 1)$. A pricing rule is *liquidity insensitive* if

$$p_i(\mathbf{q} + \alpha \mathbf{1}) = p_i(\mathbf{q})$$

for all valid \mathbf{q} and all α .

Sensitivity to liquidity is desirable because it aligns intuitively with the way one would want markets to function: a fixed-size investment moves prices less in thick (liquid) markets than in thin (illiquid) markets.

One can also think about sensitivity from a Bayesian perspective. The 1000th flip of a coin moves the posterior estimate of that coin’s probability of coming up heads much less than the first flip. This is because, after 1000 flips, we already have a great deal of information about the probability of the coin coming up heads. Similarly, if we have a lot of information about the objective price of a contract (a deep market), small bets in the market should not impact prices much.

2.1.2. Tension among the Desired Properties. In this section, we show that no market maker can satisfy all three of the desired properties.

Definition 2.8. Any market maker that satisfies translation invariance and path independence is a *Hanson market maker*.

This name is inspired by Robin Hanson, who provided an approach to building such market makers from strictly proper scoring rules. All of the example market makers given by Hanson and subsequent authors [Agrawal et al. 2009; Chen and Pennock 2007, 2010; Chen and Vaughan 2010; Pennock and Sami 2007; Peters et al. 2007] are liquidity insensitive. We now show why: liquidity sensitivity is in fact impossible to achieve in the Hanson context.

THEOREM 2.9. *No pricing rule is translation invariant, path independent, and liquidity sensitive.*

PROOF. We show that a Hanson market maker, which is by definition translation invariant and path independent, has constant prices along $\mathbf{1}$ and is therefore liquidity insensitive.

Because Hanson market makers are path independent, prices are given by the gradient of a scalar field, the cost function. Consider the Hessian of that cost function

$$\nabla^2 C(\cdot) = \begin{bmatrix} \frac{\partial p_1}{\partial q_1} & \dots & \frac{\partial p_n}{\partial q_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial p_1}{\partial q_n} & \dots & \frac{\partial p_n}{\partial q_n} \end{bmatrix}.$$

The sum of the entries of the i th row of this matrix represents the change in the sum of prices from adjusting q_i . Since prices always sum to 1, each row of the matrix sums to 0.

By the symmetry of second derivatives, the columns of the Hessian are identical to the rows of the Hessian. Therefore, each column of the Hessian also sums to 0. So we have

$$\sum_j \frac{\partial p_j}{\partial q_i} = 0 = \sum_i \frac{\partial p_j}{\partial q_i} = \mathbf{1} \cdot \nabla p_j = \nabla_{\mathbf{1}} p_j,$$

where $\nabla_{\mathbf{1}}$ represents the directional derivative along $\mathbf{1}$. Since the directional derivatives of the prices along $\mathbf{1}$ are all 0, the prices are constant along $\mathbf{1}$, and so Hanson market makers are liquidity insensitive. \square

3. INTRODUCING OUR MARKET MAKER

As we have discussed, a Hanson rule satisfies translation invariance and path independence; it is not sensitive to liquidity and it will not make money in expectation. In this section, we introduce our market maker, which is path independent, adaptive to increased liquidity, and can arrive at situations in which it makes money regardless of realized outcome. We delve into the theoretical properties of our market maker in detail in Section 4.

We begin this section by considering two ways of modifying a Hanson rule in an effort to make it more practical. Though we dismiss both approaches, their rejection helps us frame the properties of the market maker we do end up constructing.

3.1. Imposing a Transaction Cost and Subsidizing Liquidity

One approach to make Hanson rules more practical is to directly impose a transaction cost on each trade. That is, bets are calculated from the Hanson rule, but an additional charge (e.g., 3%) is added to every transaction presented to a potential bettor. For instance, if we present a trader with a bet that would normally cost 1 dollar according to the Hanson rule, it would instead cost 1.03. The market maker can then keep 0.03.

Imposing a transaction cost enables a market maker to potentially run at a profit, assuming a sufficient level of market activity. However, this scheme is still not liquidity sensitive—prices respond identically to bets at all different volumes.

A second and more complex idea is to break the transaction fee between increasing liquidity and collecting a fee. To our knowledge, this idea was originally proposed by Todd Proebsting. For instance, a market maker can charge a 3% fee, but only keep 1%, putting the other 2% towards increasing liquidity (perhaps by increasing the amount of liquidity so that the worst-case loss is larger by the amount of the 2% subsidy). Such a market maker would be liquidity sensitive and can run at a profit, but has two shortcomings.

The first shortcoming is that increasing liquidity in this manner has a tendency to distort prices towards $1/n$. Recall that liquidity is a measure of how much the market maker adjusts marginal prices in response to market activity. At higher levels of liquidity, it takes larger magnitude quantity vectors to produce the same prices.

Therefore, increasing liquidity in the LMSR has the effect of dampening extreme prices by pushing prices closer together. Another way of viewing this effect is to consider the equivalence between (convex) Hanson market makers and regularized online follow-the-leader algorithms explored in Chen and Vaughan [2010]. As those authors show, the concept of liquidity in automated market makers is analogous to the amount of regularization applied to prices—that is, with more liquidity, prices are closer to an initial prior, which is generally a uniform estimate of $1/n$ over each state (this is the case for the LMSR). It is of particular concern that agents, knowing this effect, could speculatively trade with an eye towards it occurring. For example, a speculator can bet on low-probability events with the understanding that future trade in the market will increase liquidity, and therefore increase the value of these bets as their prices move toward the mean.

The second shortcoming is that it breaks path independence, because a series of smaller orders will result in more updates to the liquidity parameter (e.g., the b term in the LMSR) than a single large order.

Our market maker can be thought of as a way of adapting this scheme *continuously* with order volume, so that prices are not distorted and so that path independence is maintained.

3.2. Relaxing Translation Invariance

Since we cannot satisfy all three of our desiderata (path independence, translation invariance, and liquidity sensitivity) simultaneously, we should consider which of them to relax. As we have discussed, Hanson market makers relax liquidity sensitivity. A more reasonable desideratum to relax is translation invariance, because it does not match how we would expect a market to function in the real world. In particular, one would like a market maker to be able to derive a profit from transacting with traders. So rather than enforcing the translation invariance condition

$$\sum_i p_i(\mathbf{q}) = 1$$

for all valid \mathbf{q} , we would actually prefer

$$\sum_i p_i(\mathbf{q}) \geq 1.$$

That way, if traders cannot take on negative quantities, the prices they face always sum to at least one.

3.3. Moving Forward in Obligation Space

Of course, with a path independent market maker, if it costs more than one dollar to acquire a dollar guaranteed payout, a trader could arbitrage the market maker by selling dollar guaranteed payouts to the market maker for more than a dollar.

One way to get around this problem is to only allow the obligation space to move forward. In this section we present two closely related ways to accomplish this goal.

3.3.1. No Selling. In this scheme, traders always purchase shares on outcomes from the market maker. Formally, let the market be at state \mathbf{q}^0 , and let a trader attempt to impose an obligation \mathbf{q} on the market maker, where

$$\min_i q_i < 0.$$

Let

$$\bar{q} \equiv -\min_i q_i.$$

Under the usual cost function scheme, that trader would pay

$$C(\mathbf{q}^0 + \mathbf{q}) - C(\mathbf{q}^0)$$

but instead, in an always moving forward scheme, the trader pays

$$C(\mathbf{q}^0 + \mathbf{q} + \bar{q}\mathbf{1}) - \bar{q} - C(\mathbf{q}^0)$$

and the market maker moves to the new state

$$\mathbf{q}^0 + \mathbf{q} + \bar{q}\mathbf{1}$$

noting that the vector $\mathbf{q} + \bar{q}\mathbf{1}$ consists of all non-negative components. This is what we mean by the market maker always moving forward in obligation space.

This scheme is still fully expressive, because with an exhaustive partition over future events the logical condition of betting against an event is equivalent to the logical condition of betting for its complement. Essentially, traders can take on the same obligations as in a traditional scheme, only they will cost more. Furthermore, if

$$\sum_i p_i(\mathbf{q}) > 1,$$

then with this scheme when a trader imposes an obligation and then sells it back to the market maker, the trader ends up with a net loss—just like the markets we see in the real world.

3.3.2. Covered Short Selling. In this scheme, traders are allowed to sell back to the market maker contracts that they have previously purchased, but are not allowed to directly sell contracts to the market maker.

Let \mathbf{q}^t represent the vector of payoffs held by trader t , so that q_i^t represents the amount the market maker must pay out to trader t if the i th event occurs. In a covered short selling scheme, the cost function operates as usual unless trader t suggests a trade that would result in

$$\min_i q_i^t < 0$$

Then, similar to the no selling scheme discussed above, the trader's payoff vector is translated by $\bar{t} \equiv -\min_i q_i^t$, so that instead the trader acquires the vector

$$\mathbf{q}^t + \bar{t}\mathbf{1}$$

noting that for all events i ,

$$(\mathbf{q}^t + \bar{t}\mathbf{1})_i \geq 0$$

The operation of selling any contract previously purchased from the market maker does not result in a trader holding a negative payoff on any event. Consequently, in this scheme, traders can buy and then immediately sell back contracts from the market maker at no net cost.

3.3.3. Discussion. Even though both schemes use the same cost function, they will produce distinct market makers when used with the cost functions we develop later in this section. A market maker that allows covered short selling permits a trader to buy and then immediately sell at no net cost. With a no selling scheme, a trader that buys and then immediately sells will incur a small loss. Which scheme is better depends on the setting; if the set of traders is sophisticated and profitability is a concern, then the no selling scheme is a better choice because it weakly dominates in terms of revenue

for the same set of trades. However, if some traders are unsophisticated and user experience is a concern, then the covered short selling scheme could be a better choice because it will not punish users for mistaken bets that they quickly cancel.

In contrast, Hanson market makers operating with either scheme (or with no scheme at all) produce exactly the same quoted costs. Let H be a Hanson market maker. Then because H is translation invariant

$$H(\mathbf{q}^0 + \mathbf{q} + \bar{q}_1) - \bar{q} = H(\mathbf{q}^0 + \mathbf{q} + \bar{q}_1 - \bar{q}_1) = H(\mathbf{q}^0 + \mathbf{q}).$$

3.4. The Logarithmic Market Scoring Rule (LMSR)

Our pricing rule is derived from the *logarithmic market scoring rule (LMSR)* [Hanson 2003]. The LMSR uses the cost function

$$C(\mathbf{q}) = b \log \left(\sum_i \exp(q_i/b) \right),$$

where $b > 0$ is the constant liquidity parameter. This function's pre-image is the entire space \mathbb{R}^n . The function's gradient, the pricing rule, is

$$p_i(\mathbf{q}) = \frac{\exp(q_i/b)}{\sum_j \exp(q_j/b)}.$$

This cost function has worst-case loss $b \log n$ for the market maker. (This loss is achieved by starting from identical prices on all events.)

3.5. Defining our Market Maker

The conventional LMSR cost function can be written as

$$C(\mathbf{q}) = b(\mathbf{q}) \log \left(\sum_i \exp(q_i/b(\mathbf{q})) \right),$$

where $b(\mathbf{q}) = b$ is an exogenously set constant. Instead, our market maker uses the LMSR cost function, but with a variable $b(\mathbf{q})$ that increases with market volume as follows:

$$b(\mathbf{q}) = \alpha \sum_i q_i,$$

where $\alpha > 0$ is a constant. The valid region for our market maker is the set of n -dimensional vectors with all non-negative components (i.e., the positive orthant), omitting the origin. In order to stay in this region, we always move forward in obligation space, as described in Section 3.3.

While it is straightforward that our market maker is path independent (because it has a cost function), it remains to be shown that it is liquidity sensitive, or in a larger sense, has any desirable qualities at all. In the next section, we explore the properties of our market maker in depth.

4. PROPERTIES OF OUR MARKET MAKER

Even though our modification to the LMSR is simple, it results in a cascade of intriguing properties.




Fig. 1. In a 2-event market with $\alpha = .05$, this plot illustrates the relationship between q_x and p_x for $q_y = 250, 500, \text{ and } 750$, respectively. The liquidity sensitivity of our market maker is evident in the decreasing slope of the price response for increasing q_y .

4.1. Prices

In a path-independent market maker, the price of state i is given by the partial derivative of the cost function along i . With constant b , this expression is simply

$$p_i(\mathbf{q}) = \frac{\exp(q_i/b)}{\sum_j \exp(q_j/b)}$$

When $b(\mathbf{q}) = \alpha \sum_i q_i$, however, the expression becomes more complex, but still analytically expressible:

$$p_i(\mathbf{q}) = \alpha \log \left(\sum_j \exp(q_j/b(\mathbf{q})) \right) + \frac{\sum_j q_j \exp(q_i/b(\mathbf{q})) - \sum_j q_j \exp(q_j/b(\mathbf{q}))}{\sum_j q_j \sum_j \exp(q_j/b(\mathbf{q}))}$$

Figure 1 illustrates the liquidity sensitivity of these prices in a 2-event market. As the number of shares of the complementary event increases, the market’s price response for a fixed-size investment becomes less pronounced.

Figures 2 and 3 show the price of a one-unit bet at various levels of liquidity in a two-event market. Figure 2 shows the price of a one-unit bet when the two events have equal quantities outstanding, while Figure 3 has the first event with proportionately higher quantities outstanding. Thus, the unit bet is more expensive in the former than the latter. Though the two figures differ quantitatively, they agree qualitatively: the price of a fixed-size contract shrinks as the level of outstanding quantities increase.

Figures 2 and 3 also illustrate an important distinction in our market maker between *instantaneous prices* and *cumulative prices*. Even though, as we show in the next section, the sum of instantaneous prices (i.e., the marginal price for a vanishingly small quantity) is bounded quite modestly for all possible outstanding quantities, at low levels of liquidity these instantaneous prices increase quite quickly. Thus, at very

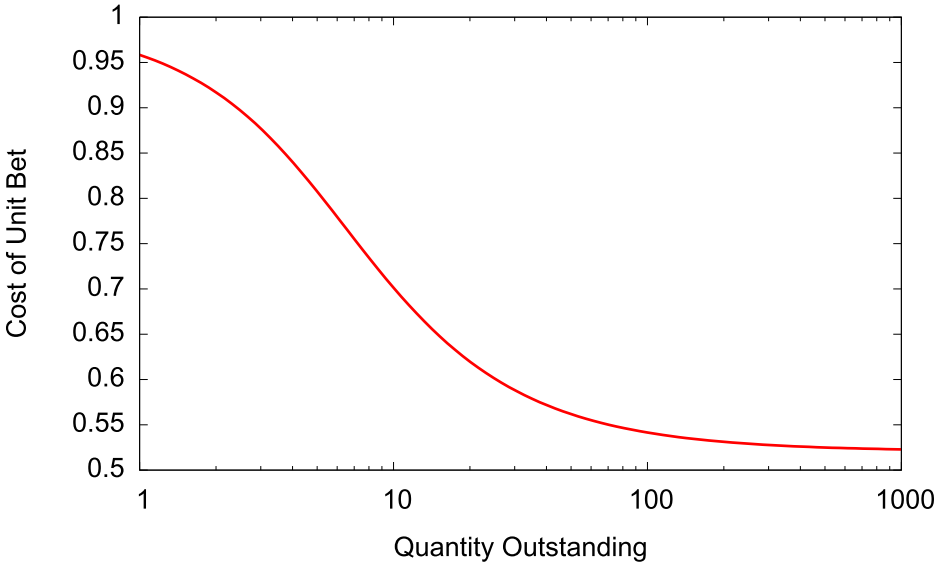


Fig. 2. In a 2-event market with $\alpha = .05$, this plot illustrates the cost of a unit bet on the first outcome when both outcomes have the designated outstanding quantity.

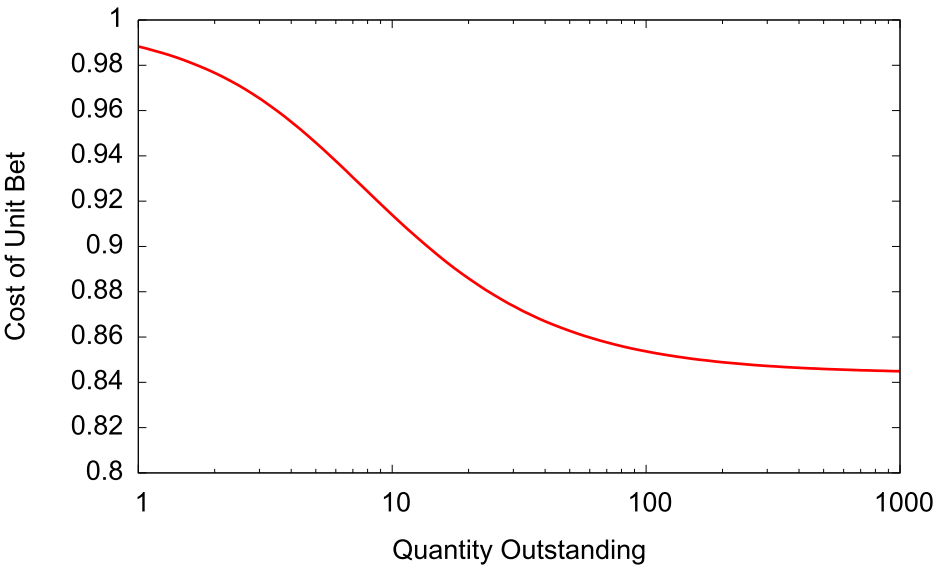


Fig. 3. In a 2-event market with $\alpha = .05$, this plot illustrates the cost of a unit bet on the first outcome when the first outcome has ten percent greater quantity outstanding than the second outcome, where the second outcome's quantity is listed (i.e., a value of 10 corresponds to (11, 10)).

small outstanding quantities, the cost of a unit bet is more than 90 cents, because our market maker is very sensitive to bets of large size relative to the quantities outstanding. At higher levels of outstanding quantities, an additional unit bet is relatively small and cumulative prices do not increase much past instantaneous prices.

4.2. Tight Bounds on the Sum of Prices

In this section, we establish tight bounds on the sum of prices. In particular, we show that

$$1 \approx 1 + n \left[\alpha \log(\exp(1/\alpha) + n - 1) - \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} \right] \leq \sum_i p_i(\mathbf{q}) \leq 1 + \alpha n \log n$$

and therefore our market maker achieves the desirable liquidity-sensitivity properties we discussed in Section 3.2.

Prices achieve their upper bound only when $\mathbf{q} = k\mathbf{1}$ for $k > 0$. Recall that $\mathbf{1}$ is the vector where each element is a 1, so the product $k\mathbf{1}$ yields a vector where each element is a k . Prices achieve the lower bound as $q_i \rightarrow \infty$.

PROPOSITION 4.1. *Prices at $k\mathbf{1}$, for all $k > 0$, sum to $1 + \alpha n \log n$.*

PROOF. For $\mathbf{q} = k\mathbf{1}$, we have $q_i = q_j$ for all i and j , which allows us to simplify considerably.

$$\begin{aligned} \sum_i p_i(k\mathbf{1}) &= \sum_i \alpha \log \left(\sum_j \exp(q_j/b(\mathbf{q})) \right) \\ &= n\alpha \log \left(\sum_j \exp(q_j/b(\mathbf{q})) \right) \\ &= n\alpha \log \left(n \exp \left(\frac{1}{\alpha n} \right) \right) \\ &= n\alpha \log \left(\exp \left(\frac{1}{\alpha n} \right) \right) + n\alpha \log n \\ &= 1 + \alpha n \log n. \quad \square \end{aligned}$$

PROPOSITION 4.2. *The maximum of the sum of prices is obtained at every point of the form $k\mathbf{1}$, where $k > 0$. Furthermore, these are the only points that achieve the maximum.*

PROOF. Consider the set of all quantity vectors that sum to $\bar{b} > 0$. We will show that the quantity vector where each event has equal quantity (each one having \bar{b}/n) maximizes the sum of prices.

The sum of prices at quantity vector \mathbf{q} is given by

$$\sum_i p_i(\mathbf{q})$$

Without loss of generality, take $\sum_i q_i = 1/\alpha$, so that the space of vectors we consider are those for which $b(\mathbf{q}) = 1$.

So, without loss of generality, we can rewrite the sum of prices as

$$1 + n\alpha \left[\log \left(\sum_j \exp(q_j) \right) - \frac{\sum_j q_j \exp(q_j)}{\sum_j \exp(q_j)} \right]$$

We will show that

$$\log \left(\sum_j \exp(q_j) \right) - \frac{\sum_j q_j \exp(q_j)}{\sum_j \exp(q_j)} \leq \log n,$$

with equality occurring only when $\mathbf{q} = k\mathbf{1}$. We can rewrite the above expression as

$$\sum_j q_j \exp(q_j) \geq \left(\sum_j \exp(q_j) \right) \log \left(\frac{\sum_j \exp(q_j)}{n} \right)$$

Take $p_j \equiv \exp(q_j)$. The expression then becomes

$$\sum_j p_j \log(p_j) \geq \sum_j p_j \log \left(\frac{\sum_j p_j}{n} \right)$$

Without loss of generality, we can scale the p_j to define a probability distribution, to get

$$\begin{aligned} \sum_j p_j \log(p_j) &\geq \log \left(\frac{\sum_j p_j}{n} \right) \\ &\geq -\log(n) \end{aligned}$$

This is a result from basic information theory, which establishes that the uniform distribution has maximum entropy over all possible probability distributions [Cover and Thomas 2006]. Therefore, equality holds only in the case of a uniform distribution, which corresponds to the quantity vector having equal components ($\mathbf{q} = k\mathbf{1}$). \square

PROPOSITION 4.3. *At any valid \mathbf{q} , $\sum_i p_i(\mathbf{q}) \geq 1$.*

PROOF. Define

$$r_i \equiv \frac{q_i}{b(\mathbf{q})}$$

and

$$s_i \equiv \frac{\exp(r_i)}{\sum_j \exp(r_j)}$$

Observe that the s_i form a probability distribution. Then, using the entropy operator H :

$$H(\mathbf{s}) = - \sum_i x_i \log x_i$$

we can express prices as

$$p_i(\mathbf{q}) = s_i + \alpha H(\mathbf{s})$$

and therefore the sum of prices as

$$\sum_i p_i(\mathbf{q}) = 1 + \alpha n H(\mathbf{s}) \geq 1.$$

Because the entropy operator is bounded below by zero, the sum of prices is at least 1. \square

There are two ways to produce a zero entropy distribution of the s_i in the above result.

- Were our market maker defined over all of \mathbb{R}^n , we could produce a zero entropy distribution by sending $q_i \rightarrow \infty$ and $q_j \rightarrow -\infty$ for $i \neq j$. However, our market maker is not defined over all of \mathbb{R}^n , but rather only in the positive orthant.

—As $\alpha \downarrow 0$, the entropy of the distribution of the s_i can approach 0. Letting q_i be positive and $q_j = 0$ for $j \neq i$, we have

$$r_i = 1/\alpha \quad \text{and} \quad r_j = 0$$

and therefore

$$s_i = \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} \quad s_j = \frac{1}{\exp(1/\alpha) + n - 1}$$

a distribution which, for fixed n , approaches a unit mass on s_i as $\alpha \downarrow 0$.

Consequently, for fixed positive α , the distribution of the s_i can have nearly zero entropy, but cannot achieve absolutely zero entropy. Thus the minimum sum of prices is not unity but rather very close to it, equal to unity to first order and well within machine precision for small values of α . The following proof formalizes this.

PROPOSITION 4.4. *The minimum sum of prices is*

$$1 + n \left[\alpha \log(\exp(1/\alpha) + n - 1) - \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} \right].$$

This minimum is achieved when $q_i > 0$ and $q_j = 0$ for $i \neq j$. For small $\alpha \gtrsim 0$,

$$1 + n \left[\alpha \log(\exp(1/\alpha) + n - 1) - \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} \right] = 1 + O(\alpha^2).$$

PROOF. From our result, the minimum sum of prices is achieved when the distribution of the s_i has minimum entropy. When restricted to the positive orthant, the corresponding distribution with largest entropy sets one q_i to be positive and the other $q_j = 0$ where $j \neq i$.

At these values, we have

$$p_i(\mathbf{q}) = \alpha \log(\exp(1/\alpha) + n - 1)$$

and

$$p_j(\mathbf{q}) = \alpha \log(\exp(1/\alpha) + n - 1) + \frac{1 - \exp(1/\alpha)}{\exp(1/\alpha) + n - 1}.$$

Observe that $p_i \approx 1$ and $p_j \approx 0$.

Adding these terms together and simplifying we get that the sum of prices is

$$1 + n \left[\alpha \log(\exp(1/\alpha) + n - 1) - \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} \right].$$

Within the braces, the left term is larger than unity while the right term is smaller than unity, meaning that the sum of prices as a whole is greater than unity, which is to be expected from our previous result.

As we will discuss, it is natural for α to be set very small. Let

$$f(\alpha) = n \left(\alpha \log(\exp(1/\alpha) + n - 1) - \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} \right).$$

Then the Taylor series of the sum of prices on the axes, taken around $\alpha = 0$, is given by

$$\sum_i p_i = 1 + f(0) + \alpha f'(0) + O(\alpha^2).$$

Since

$$\lim_{\alpha \downarrow 0} \alpha \log(\exp(1/\alpha) + n - 1) - \frac{\exp(1/\alpha)}{\exp(1/\alpha) + n - 1} = 1 - 1 = 0,$$




Fig. 4. Sum of prices where $n = 2$ and $\alpha = 0.05$. The sum is bounded between 1 and $1 + \alpha n \log n \approx 1.07$, achieving its maximum where $q_x = q_y$.

the $f(0)$ term of the expression is zero, meaning that the total deviation away from 1 for small α is given by the term $\alpha f'(0)$. The derivative is a complex expression that we give for completeness:

$$f'(\alpha) = n \left(\frac{e^{1/\alpha}}{\alpha^2(n + e^{1/\alpha} - 1)} - \frac{e^{2/\alpha}}{\alpha^2(n + e^{1/\alpha} - 1)^2} - \frac{e^{1/\alpha}}{\alpha(n + e^{1/\alpha} - 1)} + \log(n + e^{1/\alpha} - 1) \right).$$

By taking the limit of this expression, we see that

$$\lim_{\alpha \downarrow 0} f'(\alpha) = 0.$$

Thus, for small α , the sum of prices is bounded below by

$$1 + O(\alpha^2)$$

Put another way, to first order the lower bound of the sum of prices of our market maker is 1. \square

Figure 4 is a plot of the sum of prices in a simple two-quantity market. Prices achieve their highest sum when $q_x = q_y$ and are bounded below by 1.

4.3. Selecting α

A possible complaint about our scheme is that we have replaced one *a priori* fixed value, b , of the LMSR with another *a priori* fixed value, our α . In this section, we discuss how the α parameter has a natural interpretation that makes its selection relatively straightforward.

The α parameter can be thought of as the commission taken by the market maker. Higher values of α correspond to larger commissions, which leads to more revenue. At the same time, setting α too large discourages trade.

As we have shown, the sum of prices with our market maker is bounded by $1 + \alpha n \log n$, and this value is achieved only when all quantities are equal. This bound provides a guide to help set α .

How large should administrators set α within our market maker? We can look to existing market makers (and bookies) for an answer. Market makers generally operate with a commission of somewhere between 2 and 20 percent. To emulate a commission that does not exceed v in our market maker, one can simply set

$$\alpha = \frac{v}{n \log n}.$$

So, the larger the event space (larger n), the smaller α should be set to maintain a given percentage commission.

Though the sum of prices increases in α , this provides no guidance as to the behavior of the cost function itself—it is not immediate that the cost function increases in α , because it has conflicting effects within our cost function. Increasing α decreases the terms $q_i/b(\mathbf{q})$ in the cost function, but scales up the output of the log function. However, the following proposition establishes that our cost function is non-decreasing in α . We are assisted in this result by the following lemma.

LEMMA 4.5. *For our cost function*

$$C(\mathbf{q}) \geq \max_i q_i.$$

PROOF. Suppose there exists a valid \mathbf{q} such that

$$C(\mathbf{q}) < \max_i q_i$$

without loss of generality, let

$$q_1 = \max_i q_i$$

and define

$$r_i = \frac{q_i}{b(\mathbf{q})} \geq 0,$$

then we have

$$\begin{aligned} \log \left(\sum_i \exp(r_i) \right) &< r_1 \\ \sum_i \exp(r_i) &< \exp(r_1) \\ \sum_{i \neq 1} \exp(r_i) &< 0, \end{aligned}$$

which is a contradiction because $\exp(x)$ is non-negative for all x . □

PROPOSITION 4.6. *Our cost function is nondecreasing in α .*

PROOF. This result follows if we can show

$$\frac{\partial}{\partial \alpha} C(\mathbf{q}) \geq 0.$$

After taking the partial derivative of our cost function and simplifying, we get

$$\left(\sum_i \exp(q_i/b(\mathbf{q})) \right) C(\mathbf{q}) \geq \sum_i q_i \exp(q_i/b(\mathbf{q})).$$

From Lemma 4.5, we have

$$C(\mathbf{q}) \geq \max_i q_i$$

and so

$$\left(\sum_i \exp(q_i/b(\mathbf{q})) \right) C(\mathbf{q}) \geq \left(\sum_i \exp(q_i/b(\mathbf{q})) \right) \left(\max_i q_i \right) \geq \sum_i q_i \exp(q_i/b(\mathbf{q})),$$

which completes the proof. \square

Recalling that the cost function defines the amount paid into the market maker, an informal way to interpret this result is that the market maker's revenue increases with the α parameter for any given quantity vector. Of course, increasing α results in higher prices, which can affect trader behavior, so the overall effect in practice might be ambiguous.

4.4. Bounded Loss

Like the LMSR, our market maker has bounded loss. In this section, we first explore why having some possible loss is actually desirable for a market maker, and then prove that our market maker has finite, arbitrarily small loss.

When pricing an obligation, a market maker could price it at least as high as the payout a trader would receive in every state of the world. But then it would not be rational for any trader to accept these offered bets. For it to be rational for a trader to accept a bet with the market maker, the bets the market maker offers must therefore (at least sometimes and possibly always) expose the market maker to a worst-case loss.

On the other hand, it is highly undesirable for a market maker to lose an infinite amount in some cases—particularly if we are using real money.

Definition 4.7. The loss of a market maker that starts in state \mathbf{q}^0 and ends in state \mathbf{q} , with the realization of event i , is

$$C(\mathbf{q}^0) - C(\mathbf{q}) + q_i$$

Recall that here $C(\cdot)$ is the cost function and q_i is the amount the market maker has to pay out in the end upon event i occurring.

Definition 4.8. A pricing rule has *bounded loss* if for all initial states \mathbf{q}^0 and all states \mathbf{q} ,

$$C(\mathbf{q}^0) - C(\mathbf{q}) + \max_i q_i < \infty$$

PROPOSITION 4.9. *Our pricing rule has bounded loss. Specifically, its loss is bounded by $C(\mathbf{q}^0)$.*

PROOF. By Lemma 4.5

$$C(\mathbf{q}) \geq \max_i q_i$$

and so

$$\max_i q_i - C(\mathbf{q}) \leq 0 \Rightarrow C(\mathbf{q}^0) + \max_i q_i - C(\mathbf{q}) \leq C(\mathbf{q}^0)$$

so our market maker's loss is bounded by

$$C(\mathbf{q}^0). \quad \square$$

Since

$$\lim_{\mathbf{q} \rightarrow \mathbf{0}} C(\mathbf{q}) = 0,$$

setting the initial market quantities close to $\mathbf{0}$, the worst-case loss becomes arbitrarily small. But reducing the initial vector too much comes at a cost, however, because

$$\lim_{\mathbf{q} \rightarrow \mathbf{0}} b(\mathbf{q}) = 0$$

so the market becomes arbitrarily sensitive to small bets in its initial stage.

In contrast, to get near-zero loss in the LMSR, one would have to set b near zero, which would cause arbitrary sensitivity to small bets throughout the duration of the market. Since other Hanson market makers are not liquidity sensitive either, they suffer from the same problem. In our market maker, by setting the initial quantities close to zero, we achieve near-zero loss while containing the high sensitivity to the initial stage only.

4.5. Worst-Case Revenue

In addition to always having bounded loss (and near-zero loss if desired), under broad conditions on the final quantity vector of the market, we can guarantee that our market maker actually makes a profit (regardless of which event gets realized). The worst-case revenue is

$$\mathfrak{R}(\mathbf{q}) \equiv C(\mathbf{q}) - \max_i q_i - C(\mathbf{q}^0).$$

If $\mathfrak{R}(\mathbf{q}) > 0$ when the market closes, the market maker will book a profit regardless of the outcome that is realized. We say that in such states the market maker has *outcome-independent profit*. Figures 5 and 6 show the set of market states for which $\mathfrak{R}(\mathbf{q}) > 0$ for various values of α and initial quantity vectors \mathbf{q}^0 .

Figure 5 shows varying values of α . From Theorem 4.6, the cost function is non-decreasing in α , which is reflected by the increasing areas of outcome-independent profit as α gets larger. Figure 6 shows varying initial quantity vectors. Since revenue is trivially decreasing in the cost of the initial quantity vector, as the cost of our initial quantity vector increases, the area of outcome-independent profit shrinks.

From the figures, it might appear that large portions of the state space will result in our market maker losing money. However, prices and quantities have a highly non-linear relationship: prices quickly approach 1 as quantities become imbalanced. The straight black rays on the plane represent a price of .95 for one of the two events. Therefore, the plots indicate that as long as markets are terminated while events have reasonable levels of uncertainty (i.e., where the price of one event is not asymptotically close to unity), the market maker can book a profit regardless of the realized future.

Figure 7 contrasts the revenue of our market maker against the LMSR. In particular, the figure shows the revenue *surplus* of the LMSR relative to our market maker. Positive values represent how much more our market maker would collect if the market




Fig. 5. The shaded regions show where the market maker has outcome-independent profit in a two-outcome market with initial quantity vector (1, 1) and various values of α . Figure (a) sets α equal to .01, Figure (b) equal to .03, and Figure (c) equal to .06. The top black ray represents $p_y = .95$ and the bottom black ray represents $p_x = .95$.




Fig. 6. The shaded regions show where the market maker has outcome-independent profit in a two-outcome market with $\alpha = .03$ and various initial quantity vectors. Figure (a) sets \mathbf{q}^0 equal to (.5, .5), Figure (b) equal to (1, 1), and Figure (c) equal to (2, 2). The top black ray represents $p_y = .95$ and the bottom black ray represents $p_x = .95$.

terminates in the each obligation state. The comparison between the two market makers is valid because both the market makers have the same bound on worst-case loss, set by aligning the α and \mathbf{q}^0 parameters in our market maker with the b parameter in the LMSR. What is especially notable is how large the revenue difference between the two market makers becomes for lopsided obligation vectors, when the market maker has to pay out much more if one event happens than if the other event happens. As Figures 5 and 6 showed, generally at lopsided obligation vectors our market maker does not book an outcome-independent profit. However, as Figure 7 shows, our market maker delivers significantly less loss than the LMSR for lopsided obligation vectors.

4.6. Homogeneity

Recall that a positive homogeneous function f of degree k has

$$f(\gamma x) = \gamma^k f(x)$$

for $\gamma > 0$. “Positive homogeneous functions of degree one” are often referred to as just “positive homogeneous”. As it turns out, the cost function of our market maker is positive homogeneous, and in this section we prove and explore the implications of that result.

PROPOSITION 4.10. *Our cost function is positive homogeneous of degree one.*




Fig. 7. A plot of the revenue surplus between the our market maker and the LMSR. The z-axis is how much more our market maker makes than the LMSR. The parameters are aligned so that the two market makers have the same worst-case loss (~ 104.2), reflected by the zero revenue surplus at \mathbf{q}^0 . In our market maker, $\alpha = .03$ and $\mathbf{q}^0 = (100, 100)$, and in the LMSR, $b = 150.27$.

PROOF. Let $\gamma > 0$ be a scalar and \mathbf{q} be some valid quantity vector. Without loss of generality, we can assume $\sum_i q_i = 1$. Then,

$$\begin{aligned}
 C(\gamma \mathbf{q}) &= b(\gamma \mathbf{q}) \log \left(\sum_i \exp(\gamma q_i / b(\gamma \mathbf{q})) \right) \\
 &= \gamma \alpha \log \left(\sum_i \exp \left(\frac{\gamma q_i}{\gamma \alpha} \right) \right) \\
 &= \gamma C(\mathbf{q}) \quad \square
 \end{aligned}$$

It is crucial that the cost function be positive homogeneous, because that allows the price response to scale appropriately in response to increased quantities. One of the primary concerns about using the LMSR is the relation of the fraction of wealth invested in the market to the displayed prices. If the b parameter is set too low in the LMSR, that is, if the market is thick but the market maker’s price response is too sensitive, then tiny fractions of the overall wealth in the market can move prices a great deal. On the other hand, if the b parameter is set too high all the wealth in the market would be insufficient to move prices significantly enough to reflect this skewed distribution of bets.

A market maker would ideally provide a price response proportional to the amount of wealth in the market. Such a market maker would appropriately scale liquidity, requiring progressively larger trades to achieve the same price response as the market accumulated more and more money. Scaling price responses proportional to the state of the market is the correct liquidity-sensitive behavior because it yields a relative price response that is the same regardless of whether the amount of money in the market is tens, thousands, or millions of dollars. Another way of thinking about this

property is that a proportional-scaling market maker is currency independent: without any further adjustment it will function equally as well regardless of whether trading is done in millions of yen or fractions of a dollar, because only the relative, rather than absolute, amounts wagered affect the market maker's price response. This leads us to the following definition.

Definition 4.11. Prices scale proportionately if

$$p_i(\mathbf{q}) = p_i(\gamma \mathbf{q})$$

for all i , \mathbf{q} and scalar $\gamma > 0$.

In fact, only homogeneous cost functions provide this price response.

PROPOSITION 4.12. *Prices scale proportionately if and only if the cost function is positive homogeneous of degree one.*

PROOF. Proportional scaling is equivalent to the price functions being positive homogeneous of degree zero. Since the k th derivative of a positive homogeneous function of degree d is itself a positive homogeneous function of degree $d - k$, if and only if the cost function is positive homogeneous of degree one will prices scale proportionately. \square

5. DISCUSSION

Two of the main practical hurdles to more widespread use of Hanson's LMSR market maker are (1) the liquidity level b is set manually and never changes, and (2) the operator can expect to lose money in proportion to b . We presented a new automated market maker design that overcomes both of these hurdles while retaining path independence, thus ensuring the market maker cannot be exploited and greatly simplifying the implementation. We proved that if we want sensitivity to liquidity and path independence, then we must relax the translation invariance condition that constrains prices of disjoint and exhaustive assets to sum to exactly one dollar. In our case, prices can sum to more than one, meaning that they no longer directly map into probabilities. However, as we showed, prices summing to more than one can enable the market maker to extract a profit from running the market if the entropy of final prices is sufficiently high. For a broad range of terminal market states our market maker actually makes a profit regardless of the event that gets realized. With the LMSR, the market operator must ante a larger subsidy to obtain reasonable liquidity. With our liquidity-sensitive market maker, the subsidy can be set arbitrarily low without harming liquidity (except in the initial stage). Perhaps most importantly in practice, our market maker is able to achieve all these properties with a simple and explicit closed form. Simplicity of representation has been one of the largest factors driving the widespread adoption of the LMSR.

While we have shown that our market maker has a broad range of new and appealing properties, they come at the consequence of forfeiting the translation invariance of the LMSR. We proved that this was necessary: no cost function-based market maker can be both liquidity sensitive and translation invariant. As a practical matter, though, losing translation invariance means losing the direct correspondence between prices and probabilities that the LMSR enjoys. Instead, what we are left with is a *range* of possible probability estimates consistent with the prices from our market maker. For instance, when $\mathbf{q} = k\mathbf{1}$, any probability between $1/n - \alpha(n - 1) \log n$ and $1/n + \alpha \log n$ for each event is consistent with the market maker's prices. Put another way, a myopic trader that had beliefs in this range would not trade with the market maker. The

space of consistent probability estimates increases as the sum of prices increases; for the small α that is natural to our setting the range of prices is relatively small and dividing the price of each event by the sum of prices provides a simple, coherent way of normalizing prices into probabilities.

If the market administrator is omniscient and can precisely set the correct amount of liquidity in the LMSR, then the translation invariance of the LMSR, and the price-probability duality it implies, is an argument for selecting the LMSR over the market maker we have described. However, if the market administrator is not omniscient and would incorrectly guess the optimal level of liquidity within the market, then the market maker we have described here is able to set the correct level of liquidity endogenously, while the LMSR would be stuck with a bad liquidity level. This makes our market maker a better choice for domains where the volume of active trading is unknown in advance—a feature of many markets, and of many Internet prediction markets in particular.

There are several unsettled issues with our market maker, including how to incorporate prior information or learning into the way the market maker prices contracts. One possible direction is suggested by the approach of Das [2008] and Das and Magdon-Ismail [2009], which feature a heuristic market maker. However, that line of research uses flexible market makers that focus more on average-case performance with nonadversarial traders. It is difficult to see how to reconcile our market maker, that provides strong bounds on quantities like worst-case loss and sums of prices, into a heuristic framework. Another direction is taken by Chen et al. [2008] and Chen and Vaughan [2010], which explore the link between Hanson market makers and no-regret learning algorithms. However, as we showed here, it is necessary to break the duality of prices and probabilities in order to achieve liquidity sensitivity. Consequently, it is not immediately clear how the no-regret framework aligns with the approach of this paper.

Additionally, our market maker operates in an online setting where traders either accept or reject bets but do not have the option of setting persistent limit orders (e.g., “I want the payoff vector \mathbf{x} at a price not greater than p ”) that may be filled in the future. In our setting, the market maker is explicitly tasked with taking on surplus quantity, and therefore risking loss, which can be contrasted with recent work on limit order matching [Blum et al. 2006; Bredin et al. 2007]. Incorporating limit orders into the market maker is tricky because persistent orders can induce discontinuities and strange effects in the market maker’s price response as new orders cause existing limit orders to be executed. Agrawal et al. [2009] present a solution to this problem using a convex optimization technique to augment a market maker with the ability to handle persistent limit orders; since our market maker is also convex, it could be augmented in a similar fashion. However, we cannot implement the solution of Agrawal et al. [2009] directly because the framework explored in that paper relies on simplifications based on translation invariance that our market maker does not satisfy. It would be interesting to explore how to handle these persistent limit orders with our market maker, and how to mix sequential and batch order processing.

Finally, our new market maker is just one instance of the class of liquidity-sensitive market makers. Other liquidity-sensitive market makers can be developed that have different relations to worst-case loss, profit, and liquidity. For instance, it might be natural to have a liquidity-sensitive market that expands liquidity only up to a certain point, after which it reaches a state of terminal liquidity and is no longer sensitive to increased transaction volume. Large-cap equities intuitively seem to have reached this terminal state; a purchase of shares of IBM today and a purchase of shares of IBM a month from now are likely to face equivalent market depth, even though billions of

dollars will have transacted in the interim. We anticipate the introduction of many different liquidity-sensitive market makers, guided by our result that, in order to have liquidity sensitivity, a market maker must break the duality between prices and probabilities.

ACKNOWLEDGMENTS

We thank Jamie Brandon, Vincent Conitzer, Ivan Corwin, Sharad Goel, Alex Grubb, Todd Proebsting, and two anonymous referees for helpful discussions and suggestions.

REFERENCES

- Abernethy, J., Chen, Y., and Vaughan, J. W. 2011. An optimization-based framework for automated market-making. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*.
- Agrawal, S., Delage, E., Peters, M., Wang, Z., and Ye, Y. 2009. A unified framework for dynamic pari-mutuel information market design. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 255–264.
- Artzner, P., Delbaen, F., Eber, J., and Heath, D. 1999. Coherent measures of risk. *Math. Finan.* 9, 3, 203–228.
- Ben-Tal, A. and Teboulle, M. 2007. An old-new concept of convex risk measures: The optimized certainty equivalent. *Math. Finan.* 17, 3, 449–476.
- Berg, J., Forsythe, R., Nelson, F., and Rietz, T. 2001. Results from a dozen years of election futures markets research. In *Handbook of Experimental Economics Results*. 486–515.
- Blum, A., Sandholm, T., and Zinkevich, M. 2006. Online algorithms for market clearing. *J. ACM* 53, 5, 845–879.
- Bredin, J., Parkes, D., and Duong, Q. 2007. Chain: A dynamic double auction framework for matching patient agents. *J. Artif. Intell. Res.* 30, 1, 133–179.
- Carmona, R. 2009. *Indifference Pricing: Theory and Applications*. Princeton University Press.
- Carr, P., Geman, H., and Madan, D. 2001. Pricing and hedging in incomplete markets. *J. Finan. Econ.* 62, 1, 131–167.
- Chen, Y. and Pennock, D. M. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 49–56.
- Chen, Y. and Pennock, D. M. 2010. Designing markets for prediction. *AI Mag.* 31, 4.
- Chen, Y. and Vaughan, J. W. 2010. A new understanding of prediction markets via no-regret learning. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 189–198.
- Chen, Y., Fortnow, L., Lambert, N., Pennock, D. M., and Wortman, J. 2008. Complexity of combinatorial market makers. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 190–199.
- Cover, T. and Thomas, J. 2006. *Elements of Information Theory*. Wiley.
- Das, S. 2008. The effects of market-making on price dynamics. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 887–894.
- Das, S. and Magdon-Ismael, M. 2009. Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems 21 (NIPS)*. 361–368.
- Föllmer, H. and Schied, A. 2002. Convex measures of risk and trading constraints. *Finan. Stochast.* 6, 4, 429–447.
- Hanson, R. 2003. Combinatorial information market design. *Inf. Syst. Front.* 5, 1, 107–119.
- Hanson, R. 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. *J. Predict. Mark.* 1, 1, 1–15.
- Oliven, K. and Rietz, T. 2004. Suckers are born but markets are made: Individual rationality, arbitrage, and market efficiency on an electronic futures market. *Manage. Sci.* 50, 3, 336–351.
- Ostrovsky, M. 2009. Information aggregation in dynamic markets with strategic traders. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 253–254.
- Othman, A. and Sandholm, T. 2010a. Automated market-making in the large: The Gates Hillman prediction market. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. 367–376.
- Othman, A. and Sandholm, T. 2010b. When do markets with simple agents fail? In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 865–872.
- Othman, A. and Sandholm, T. 2011. Liquidity-sensitive market makers via homogeneous risk measures. In *Proceedings of the International Workshop on Internet and Network Economics (WINE)*.

- Pennock, D. and Sami, R. 2007. Computational aspects of prediction markets. In *Algorithmic Game Theory*, Cambridge University Press, 651–674.
- Peters, A., So, A. M.-C., and Ye, Y. 2007. Pari-mutuel markets: Mechanisms and performance. In *Proceedings of the International Workshop on Internet and Network Economics (WINE)*. 82–95.
- Williams, O. M. 2009. Regulatory oversight and recent initiatives to address risk posed by credit default swaps. Testimony Before the Subcommittee on Capital Markets, Insurance, and Government Sponsored Enterprises.

Received September 2011; revised April 2012, August 2012; accepted August 2012