

Profit-Charging Market Makers with Bounded Loss, Vanishing Bid/Ask Spreads, and Unlimited Market Depth

ABRAHAM OTHMAN, Carnegie Mellon University
TUOMAS SANDHOLM, Carnegie Mellon University

Four desiderata for automated market makers have appeared in the literature: (1) bounded loss, (2) the ability to make a profit, (3) a vanishing bid/ask spread, and (4) unlimited market depth. Intriguingly, market makers that satisfy *any three* of these desiderata have appeared in the literature. However, it was an open question as to whether a market maker can simultaneously satisfy all four because the qualities are oppositional. In this paper, we design market makers that satisfy all four. We achieve this by introducing a new, practical framework. It extends constant-utility cost functions with two separate functions that are added to the prices quoted to the trader. The *liquidity function* uses its proceeds to increase the amount of liquidity provided by the market maker. The *profit function* represents a “lockbox” of savings that is separate from the rest of the market maker’s decision-making process.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; I.2.11 [Distributed Artificial Intelligence]: Multi-agent Systems

General Terms: Algorithms, Design, Economics

Additional Key Words and Phrases: Market Making, Automated Market Makers, Prediction Markets, Financial Markets, Algorithmic Finance, Market Microstructure, Wagering, Liquidity, Markets

1. INTRODUCTION

Automated market makers are algorithmic agents that provide liquidity in (electronic) markets. In many markets, there may not be enough organic liquidity to support active trade, or the market may encompass enough events that buyers and sellers have trouble finding one another. Markets mediated by automated agents have successfully predicted the openings of buildings [Othman and Sandholm 2010a], point spreads in sports matches [Goel et al. 2008], anticipated the ratings of course instructors [Chakraborty et al. 2011], etc. Automated market makers are also used by a number of companies (e.g., Inkling Markets) offering private corporate prediction markets to aggregate internal information. Automated market makers are also used in financial markets in practice and in research [Othman and Sandholm 2012]. An alternate, more theoretical view of automated market makers places them as a platonic ideal of the way prices move in response to trades being made [Ostrovsky 2009; Othman and Sandholm 2010b]. Overviews of automated market making are given by Pennock and Sami [2007] and Chen and Pennock [2010].

Real-world markets are currently populated by profitable human (or human-controlled) market makers. While effective on average, their worst-case performance is

This work was supported by NSF grants CCF-1101668, IIS-0964579, and IIS-0905390, and by the Google PhD Fellowship in Market Algorithms.

Author’s addresses: Computer Science Department, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213. {aothman, sandholm}@cs.cmu.edu

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.

EC’12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

questionable (e.g., in the recent “Flash Crash” [U.S. Commodity and Futures Trading Commission and U.S. Securities & Exchange Commission 2010]). Human-controlled market makers often withdraw from uncertain or volatile markets, yielding catastrophic consequences [MacKenzie 2006; Taleb 2007]. In contrast, some market makers from the computer science literature have well-defined performance characteristics, including bounds on loss independent of the behavior of counterparties, and need not panic in the face of uncertainty. Providing an algorithmic alternative to more fickle human-controlled market makers should be considered a long-term goal of this line of research.

One obstacle that has held this line back is that the market making agents in the literature do a poor job of matching many attributes of human market makers, particularly those related to profitability. To put it bluntly, those agents are virtually assured of losing money. Automated market making agents have been employed where the goal is information elicitation, and the market maker’s losses can be rationalized as subsidies for information revelation [Hanson 2003; Chen and Pennock 2010]. However, real markets seldom run at a loss. The agents from the literature may also be subject to the vagaries of unchangeable prior parameters, like a fixed amount of market depth, which can produce problems in practice, even when used with fake money [Othman and Sandholm 2010a], or a sizable marginal bid/ask spread, yielding ambiguity over the true probabilities implied by prices [Othman et al. 2010]. The agents of the literature have not been successfully employed in any real-money markets where profit and loss are important considerations.

In this paper, we develop market making agents that provide a principled way to extend the worst-case results of the literature to incorporate desirable, realistic features in practice. These are the first agents to include three human properties: the ability to make a profit, and (provided sufficient trading volume) unlimited market depth, and a vanishing bid/ask spread. At the same time, our agents retain the bounded loss property of the algorithmic agents in the literature. Furthermore, our agents incentivize myopic traders to directly reveal their private beliefs. In addition to these desirable theoretical properties, our new market making agents have other key practical properties. Just like in real markets, but unlike most agents in the literature, our market making agents are not path independent: a trader that buys from and then sells to our market making agents will incur a small loss. Additionally, our market making agents provide a straightforward way to incorporate the principal’s subjective belief over the future into the quoted prices, and those quoted prices can be computed efficiently and simply.

2. DESIRABLE PROPERTIES FOR A MARKET MAKER

In this section, we use a real example to demonstrate how liquidity is provisioned in real markets. We then use the insights gleaned from the example to motivate four desiderata for automated market makers. Finally, we give an overview of market makers from the literature to show that, while they can achieve every combination of three of the desiderata, no existing approach satisfies all four.

2.1. Our study of stock and prediction markets

The word *liquidity* in financial markets is burdened with several connotations. O’Hara [1995] memorably quips: “liquidity, like pornography, is easily recognized but not so easily defined”. She goes on to describe several perspectives on what liquidity means: the volume in the market; the size of the marginal bid/ask spread (i.e., the spread for the smallest possible quantity); and the degree to which large bets move prevailing prices (equivalent to the volume of bets placed near the marginal bid/ask spread, the *market depth*). In traditional markets, all the senses of the term liquidity are conflated

because these characteristics tend to accompany each other. We proceed to illustrate the conflation of these several views of liquidity using live, current markets as examples.

In real markets as the volume of trade increases, the bid/ask spread falls. This implies both that a fixed size bet moves the market less and that the marginal bid/ask spread decreases. Although there can be exceptions—such as the Flash Crash—this relationship holds broadly and was discussed at least as early as 1968 [Demsetz 1968].

To provide an example of the effect of volume on bid/ask spreads, we collected values from five live markets. Table I shows the total cost to buy and then sell a thousand dollars of underlying contracts (without regard to any trading fees imposed by the exchange or brokers). We took these values from a snapshot of the relevant order books at 3pm on September 8th, 2011.

Table I. The difference in bid and ask prices for \$1,000 dollars of the underlying in several markets.

Market	Difference (\$)	Approximate Daily Volume (\$)
NWS stock	0.61	140 million
NYT stock	1.51	18 million
MNI stock	14.41	1.2 million
Obama 2012	31.58	900
Higgs Boson 2011	970	1.20

Note: “Obama 2012” is the Barack Obama re-elected contract on Intrade, and “Higgs Boson 2011” is the discovery of the Higgs Boson in 2011 contract on Intrade.

To illustrate how we calculated these values, imagine a certain security had the following order book: bid orders for 200 shares at both 2 and 3 dollars, and ask orders for 250 shares at 4 dollars. Then the bid/ask spread would be calculated as 300 dollars: 250 shares would be purchased at 4 dollars, exhausting the thousand dollar budget. Then, 200 shares would be resold at 3 dollars for a 1 dollar loss each, and the remaining 50 shares would be re-sold at 2 dollars for a 2 dollar loss each.

NWS, NYT, and MNI are all equities in news companies. The Obama 2012 Intrade contract, which pays off if the President is re-elected, is one of the most popular contracts on what is probably the most popular Internet prediction market. The Higgs Boson 2011 contract pays out if the Higgs Boson is discovered before the end of 2011 and is a very sparsely traded Intrade contract. The NWS contract is very liquid—well over a hundred million dollars worth is exchanged each day and sizable positions can be exited at almost no spread (the bid/ask spread on each share of NWS is one cent, the smallest possible value). In contrast, the Higgs Boson contract is very illiquid, seeing almost no trade each day. The bid/ask spread on a sizable position is large: buying and then selling a 1,000-dollar position results in an immediate loss of 970 dollars.

2.2. Four oppositional desiderata in the literature

The observations of the previous section suggest that real markets have a shrinking bid/ask spread for fixed-size bets as the volume gets large. This, combined with the profit motives of real market makers and the algorithmic worst-case guarantees of the market makers from the literature, yields four desiderata for market making agents:

- (1) Bounded worst-case loss.
- (2) The ability to enter a state where the market maker books a profit regardless of which future state of the world is realized.
- (3) A marginal bid/ask spread that approaches zero in the limit as volume gets large.

- (4) The price of any fixed-size transaction approaches the marginal bid or ask price, so that there is unlimited depth in the limit as volume gets large.

In addition to their self-evident desirability, *all four combinations of exactly three of these properties already exist in the literature*. Satisfying all four characteristics is challenging because several of the qualities are oppositional; for instance, making a profit involves charging extra, but charging extra means that the bid/ask spread may not vanish. As another example, a market maker that is very deep must not move prices very much in the face of large bets, but if prices are not moved enough then worst-case loss can become unbounded. To illustrate the challenge involved in satisfying all four of these desiderata, we proceed to provide a quick survey of the relevant constructions from the literature. (The next sections will define the desiderata formally.)

2.2.1. Fixed prices. Probably the simplest automated market maker is to determine a probability distribution over the future states of the world, and to offer to make bets directly at those odds. This scheme offers unlimited depth and no marginal bid/ask spread, but has unbounded worst-case loss and no ability to book a profit.

2.2.2. Fixed prices with profit-taking. Adding a profit cut on top of the fixed odds gives the market maker the ability to make a profit and retains the unlimited depth of the fixed pricing scheme. However, this market maker has unbounded worst-case loss and a non-vanishing bid/ask spread.

2.2.3. Fixed prices with shrinking profit cuts. If we allow the profit cut to diminish to zero as trading volume increases, the resulting market maker has three of the four desired properties: the ability to make a profit, a vanishing marginal bid/ask spread, and unbounded depth in limit. However, it still has unbounded worst-case loss because a trader with knowledge of the true future could make an arbitrarily large winning bet with the market maker.

2.2.4. Convex risk measures. Convex risk measures are the general class of market makers [Agrawal et al. 2009; Othman and Sandholm 2011b] featured in much of the prediction market literature [Chen and Pennock 2007; Peters et al. 2007; Agrawal et al. 2009; Chen and Vaughan 2010; Othman and Sandholm 2010a; Abernethy et al. 2011], including the most widely-used automated market maker in practice, the *Logarithmic Market Scoring Rule (LMSR)* [Hanson 2003, 2007]. These market makers can offer bounded worst-case loss and no marginal bid/ask spread. However, they do not offer the ability for the market maker to book a profit, and they offer a fixed market depth that does not increase with volume. (For instance, in the LMSR, the depth of the market is fixed by the parameter b which is an exogenous constant set *a priori*.)

2.2.5. Convex risk measures with profit-taking. Adding a fixed charge on top of a convex risk measure gives the market maker the ability to make a profit, but the marginal bid/ask spread will not vanish.

2.2.6. Convex risk measures with shrinking profit cuts. If we allow the profit cut to shrink to zero as trading volume increases, the resulting market maker can have bounded loss, the ability to make a profit, and a marginal bid/ask spread that vanishes. However, the depth of the market is still an exogenous constant, and a fixed size bet will always diverge in price from the marginal.

2.2.7. Extended constant-utility cost functions. Othman and Sandholm [2011a] describe a market maker that has unbounded depth in the limit, bounded loss, and vanishing bid/ask spread in the limit. However, it has no ability to book a profit. The market making agents of Othman and Sandholm [2011a] are a restricted special case of the

market makers we develop in this work, where the utility function and liquidity function are both logarithmic functions and there is no profit function.

2.2.8. Liquidity-sensitive automated market makers. The market maker described in Othman et al. [2010] has three of the desired properties. While it does have unbounded depth in the limit, bounded loss, and the ability to make a profit, the marginal bid/ask spread never goes to zero.

3. TECHNICAL PRELIMINARIES

This section provides a formal introduction to the automated market maker setting, which we will extend to create our market makers that satisfy all four desiderata. As is standard, the world is exhaustively partitioned into finite n possible future states $\omega_i \in \Omega$. The market maker keeps track of its state through a payout vector.

Definition 1. A *payout vector* is a vector $\mathbf{x} \in \mathbb{R}^n$ where x_i is the cumulative amount the market maker must pay out to traders if future state of the world ω_i is realized.

Definition 2. Let $U \subset \mathbb{R}$ be an open interval on the real line that includes all positive values. A *utility function* is a strictly increasing, concave function $u : U \mapsto \mathbb{R}$.

We are particularly interested in a special class of utility functions called *barrier utility functions*. These are functions which have a barrier to going into negative wealths.

Definition 3. A *barrier utility function* u is a utility function that has

$$\lim_{x \downarrow 0} u(x) = -\infty.$$

Examples of barrier utility functions include $u(x) = \log(x)$ and $u(x) = -1/x$, both of which are defined over the open interval $(0, \infty)$.

Cost function-based market makers are standard in the literature [Chen and Pennock 2007, 2010; Othman et al. 2010]. These market makers have a cost function $C : \mathbb{R}^n \mapsto \mathbb{R}$ that tracks the amount of money paid into the market maker. When a trader wants to make a bet with a market maker that would change the market maker's payout vector from \mathbf{x} to \mathbf{y} , the market maker offers the bet at a price of $C(\mathbf{y}) - C(\mathbf{x})$. Consider a two-event market where the market maker's payout vector is $(5, 3)$. If a trader wishes to place a bet that pays out a dollar if the first event happens, a cost function-based market maker would quote that trader the price of $C((6, 3)) - C((5, 3))$.

Our work extends the constant-utility cost functions originally developed by Chen and Pennock [2007].

Definition 4. Let $x^0 \in \text{dom } u$ and p_i be the market maker's (subjective) probability that ω_i will occur. A *constant-utility cost function* $C : \mathbb{R}^n \mapsto \mathbb{R}$ implicitly solves $\sum_i p_i u(C(\mathbf{x}) - x_i) = u(x^0)$.

The intuition here is that the market maker quotes prices to traders in order to maintain constant utility of $u(x^0)$. Because u is strictly increasing, costs are uniquely defined for every payout vector \mathbf{x} [Chen and Pennock 2007]. Although they are not generally expressible in closed form, constant-utility cost functions are still easy to calculate: since utility functions are strictly increasing the cost can be solved to k bits of precision in k iterations of a binary search.

One quality that is particularly important for an automated market maker to have is *bounded worst-case loss*. This gives the market administrator a performance guarantee regardless of the traders' behavior (i.e., even if a trader has perfect foresight).

Definition 5. The *worst-case loss* of a market maker is the most, over all interaction sequences and all possible ω_i , that the market maker could lose. Formally, let $\mathbb{X} = \{\mathbf{x}^1, \mathbf{x}^2, \dots\}$ be a (possibly infinite) chain of transactions, where \mathbf{x}^i denotes that the i -th trade moves the market maker from payout vector \mathbf{x}^{i-1} to payout vector \mathbf{x}^i , let $\mathcal{M}(\mathbb{X})$ represent the least upper bound on the amount paid into the market maker for the chain of transactions \mathbb{X} , and let

$$x'(\mathbb{X}) = \limsup_{k \rightarrow \infty} (\max_i x_i \in \mathbf{x}^k)$$

In this notation, the worst-case loss of a market maker is given formally by

$$\sup_{\mathbb{X}} \mathcal{M}(\mathbb{X}) - x'(\mathbb{X}).$$

(If we were to restrict our attention to path-independent market makers [Pennock and Sami 2007; Chen and Pennock 2007; Othman et al. 2010; Abernethy et al. 2011], worst-case loss would be much simpler to define, because those market makers operate only on their current payout vector, without regard to the past history of transactions.)

One advantage of using constant-utility cost functions with barrier utility functions is that their worst-case loss can be bounded simply.

PROPOSITION 1 ([OTHMAN AND SANDHOLM 2011A]). *Let C be a constant-utility cost function that employs a barrier utility function. If $p_i > 0$ for every i , the cost function-based market maker using C loses at most x^0 .*

Now that we have provided the basic framework for the automated market maker setting, we can describe how our market maker extends this framework.

4. HOW BETS ARE TAKEN

In this section, we describe how to calculate the prices an agent sees when she trades with our market makers. We denote a market maker by \mathcal{M} and we denote $\mathcal{M}(\mathbf{x}, \mathbf{y})$ to be the *total price* charged to agents for moving the market maker from payout vector \mathbf{x} to payout vector \mathbf{y} .

Let u be a utility function and $x^0 \in \text{dom } u$. Let f and g be non-decreasing functions $\mathbb{R}^+ \mapsto \mathbb{R}$ with the property that $f(0) = g(0) = 0$. For reasons that will shortly become clear, we call f the *liquidity function* and g the *profit function*. Let $d : \mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}^+$ be a distance function (metric). The state s is an internal scalar initialized so that $s = 0$. s can be thought of as a measure of the cumulative volume transacted in the market.

In order to price a bet, the following steps are taken:

- (1) A trader wishes to place a bet that would move the market maker from payout vector \mathbf{x} to payout vector \mathbf{y} . The market maker's current state is s .
- (2) The cost function $C(\mathbf{y})$ is solved implicitly for

$$\sum_i p_i u(C(\mathbf{y}) - y_i + f(s + d(\mathbf{x}, \mathbf{y}))) = u(x^0 + f(s + d(\mathbf{x}, \mathbf{y})))$$

This equation incorporates the liquidity function f but not the profit function g . Just as in constant-utility cost functions, C will be uniquely defined because u is strictly increasing, and it can be calculated efficiently through a binary search.

- (3) The total cost quoted to the trader for the bet is the sum of the changes to the cost function, liquidity function, and profit function

$$\mathcal{M}(\mathbf{x}, \mathbf{y}) \equiv C(\mathbf{y}) - C(\mathbf{x}) + f(s + d(\mathbf{x}, \mathbf{y})) - f(s) + g(s + d(\mathbf{x}, \mathbf{y})) - g(s)$$

- (4) If the bet is taken, the state changes: $s \leftarrow s + d(\mathbf{x}, \mathbf{y})$, and the new value of the cost function $C(\mathbf{y})$ is saved for the next transaction.

We will call a market making agent \mathcal{M} that prices bets this way a *constant-utility profit-charging market maker*.

The values involved in calculating \mathcal{M} all *telescope* from their initial values. Consider a transaction that first moves from payout vector \mathbf{x} to payout vector \mathbf{y} , ending at payout vector \mathbf{z} . The total amount paid into the market maker is

$$\mathcal{M}(\mathbf{x}, \mathbf{y}) + \mathcal{M}(\mathbf{y}, \mathbf{z}) = C(\mathbf{z}) - C(\mathbf{x}) + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - f(s) + g(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - g(s)$$

Observe that the terms from the intermediate payout vector \mathbf{y} (i.e., $C(\mathbf{y})$, $f(s + d(\mathbf{x}, \mathbf{y}))$, and $g(s + d(\mathbf{x}, \mathbf{y}))$) cancel out and so do not appear in this expression. We will use this telescoping property to simplify some of the proofs in Section 5.

5. THEORETICAL PROPERTIES

In this section, we show that our market makers satisfy several desirable properties. We begin by showing that they have bounded loss. As a step toward this result, we show that they have no money pumps, so they incentivize agents to directly acquire their desired portfolios (rather than taking a roundabout path). We then show that our agents have unlimited depth and no bid/ask spread in the limit. Finally, we formalize the notion of profit-taking and describe how our agents can enter a state of *unconditional profit*, in which no matter what the realized outcome is or what the future actions of the traders are the market maker will book a profit.

5.1. Bounded loss

One particularly undesirable property for a market making agent is to be in possession of a *money pump*.

Definition 6. A market maker \mathcal{M} has a *money pump* if there exists a sequence of payout vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ and some terminal state \mathbf{x}_0 such that

$$\mathcal{M}(\mathbf{x}_0, \mathbf{x}_1) + \dots + \mathcal{M}(\mathbf{x}_{n-1}, \mathbf{x}_n) + \mathcal{M}(\mathbf{x}_n, \mathbf{x}_0) < 0$$

When a market maker has a money pump a trader can keep arbitraging the market maker for unbounded riskless profit. Perhaps the most prominent difference between our work and the prior literature is that the total prices charged by our market making agents are not necessarily path independent, that is, different paths through quantity space may correspond to different costs being charged by our market making agents. With a path-independent market maker, buying a contract and then immediately selling it is without cost. This is in contrast to both real-world markets and our market making agents, where buying and then immediately selling is costly to the trader (see, e.g., Table I).

Every cost function-based market maker is path independent, and path independence is sufficient for a market maker to have no money pumps [Othman et al. 2010]. However, path independence is not *necessary* for a market maker to have no money pumps; having a path-dependent market making agent simply means that it is not immediate that the market maker has no money pumps.

Definition 7. A market maker \mathcal{M} obeys the *triangle inequality* if, for all payout vectors \mathbf{x} , \mathbf{y} , and \mathbf{z}

$$\mathcal{M}(\mathbf{x}, \mathbf{y}) + \mathcal{M}(\mathbf{y}, \mathbf{z}) \geq \mathcal{M}(\mathbf{x}, \mathbf{z})$$

By inductive argument, it is easy to see that a market maker that obeys the triangle inequality cannot have a money pump.

PROPOSITION 2. *Let \mathcal{M} be a constant-utility profit-charging market maker. Then \mathcal{M} obeys the triangle inequality.*

PROOF. Consider $\mathcal{M}(\mathbf{x}, \mathbf{y}) + \mathcal{M}(\mathbf{y}, \mathbf{z}) - \mathcal{M}(\mathbf{x}, \mathbf{z})$. In terms of the triangle inequality, this is the difference between taking the “long way” around the triangle versus the direct way; we seek to show that it is non-negative for all \mathbf{x}, \mathbf{y} , and \mathbf{z} .

Because all the values associated with the calculation of \mathcal{M} telescope, this difference is just

$$C_1(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) + g(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - (C_2(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{z})) + g(s + d(\mathbf{x}, \mathbf{z}))) \quad (1)$$

where $C_1(\mathbf{z})$ solves

$$\sum_i p_i u(C_1(\mathbf{z}) - z_i + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z}))) = u(x^0 + s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z}))$$

and C_2 solves

$$\sum_i p_i u(C_2(\mathbf{z}) - z_i + f(s + d(\mathbf{x}, \mathbf{z}))) = u(x^0 + s + d(\mathbf{x}, \mathbf{z}))$$

We will divide the terms in Equation 1 and deal with them each in turn. First, we will show that

$$C_1(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - (C_2(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{z}))) \geq 0$$

Since u is strictly increasing and d is a distance function, we have

$$u(x^0 + s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) \geq u(x^0 + s + d(\mathbf{x}, \mathbf{z}))$$

and so

$$\sum_i p_i u(C_1(\mathbf{z}) - z_i + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z}))) \geq \sum_i p_i u(C_2(\mathbf{z}) - z_i + f(s + d(\mathbf{x}, \mathbf{z})))$$

which, because u is strictly increasing and the p_i form a probability distribution, implies

$$C_1(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) \geq C_2(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{z}))$$

so

$$C_1(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - (C_2(\mathbf{z}) + f(s + d(\mathbf{x}, \mathbf{z}))) \geq 0 \quad (2)$$

Now consider the other set of terms in Equation 1,

$$g(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - g(s + d(\mathbf{x}, \mathbf{z}))$$

Because d satisfies the triangle inequality, $d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z}) \geq d(\mathbf{x}, \mathbf{z})$. Therefore, the argument to the $g(\cdot)$ function on the left is at least as large as the argument to the $g(\cdot)$ function on the right. Since $g(\cdot)$ is a non-decreasing function in its arguments

$$g(s + d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})) - g(s + d(\mathbf{x}, \mathbf{z})) \geq 0 \quad (3)$$

Putting together Equations 1, 2, and 3, we have that $\mathcal{M}(\mathbf{x}, \mathbf{y}) + \mathcal{M}(\mathbf{y}, \mathbf{z}) - \mathcal{M}(\mathbf{x}, \mathbf{z}) \geq 0$, or rewritten, $\mathcal{M}(\mathbf{x}, \mathbf{y}) + \mathcal{M}(\mathbf{y}, \mathbf{z}) \geq \mathcal{M}(\mathbf{x}, \mathbf{z})$, and so \mathcal{M} obeys the triangle inequality. ■

The proof also provides the intuition that, with an increasing profit function, an agent is charged strictly greater prices when he does not acquire inventory on the

path directly to his desired allocation. This is because the market maker will collect more money from the profit function, while not losing any additional money from the potential added liquidity from taking a longer path. Consequently, myopic agents are incentivized to directly acquire their desired allocation.

Because the total prices charged by a constant-utility profit-charging market maker obey the triangle inequality, we also have the following result.

COROLLARY 1. *Let \mathcal{M} be a constant-utility profit-charging market maker. Then \mathcal{M} has no money pumps.*

Recall that a constant-utility cost function employing a barrier utility function has a worst-case loss of the x^0 originally used to seed the utility function. Because \mathcal{M} has no money pumps, we have the following result.

COROLLARY 2. *Let u be a barrier utility function. Then if $p_i > 0$ for every i , a constant-utility profit-charging market maker \mathcal{M} that uses u loses no more than the x^0 originally used to seed the utility function, regardless of the trades made by agents or the realized future state of the world. (This result holds even if the liquidity function f and the profit function g are zero.)*

5.2. Profit, liquidity, and market depth

We begin this section by showing that, under certain conditions, a constant-utility profit-charging market maker has a vanishing bid/ask spread. Then, we show that under more stringent conditions, constant-utility profit-charging market makers also have unbounded market depth.

Throughout this section, to simplify the proofs, we assume d is a distance function implied by any \mathcal{L}_p norm. However, the results of this section hold for any continuous distance function.

5.2.1. Vanishing bid/ask spread. In this section, we show that if the profit and liquidity functions are diminishing, a constant-utility profit-charging market maker has a vanishing bid/ask spread.

Definition 8. For a market maker \mathcal{M} with a differentiable price response, let \mathcal{M}_i denote the marginal cost of a bet on the i -th event. A market maker has a *vanishing bid/ask spread* if $\sum_i \mathcal{M}_i = 1$.

PROPOSITION 3. *Let the liquidity function f and profit function g have the property that $\lim_{s \rightarrow \infty} f'(s) = \lim_{s \rightarrow \infty} g'(s) = 0$. Then a constant-utility profit charging market maker has a vanishing bid/ask spread as s gets large. Formally, $\lim_{s \rightarrow \infty} \sum_i \mathcal{M}_i = 0$.*

PROOF. At any \mathbf{x} and s , by definition we have

$$\mathcal{M}_i = \nabla_i C(\mathbf{x}) + \nabla_i f(s) + \nabla_i g(s) = \nabla_i C(\mathbf{x}) + f'(s) + g'(s)$$

Since, by construction, both $f'(s)$ and $g'(s)$ go to zero as s gets large, the interesting term is $\nabla_i C(\mathbf{x})$. From the definition of constant-utility profit-charging market makers we have

$$\nabla_i \left(\sum_j p_j u(C(\mathbf{x}) - x_j + f(s)) \right) = \nabla_i u(x^0 + f(s))$$

$$\nabla_i p_i u(C(\mathbf{x}) - x_i + f(s)) + \nabla_i \sum_{j \neq i} p_j u(C(\mathbf{x}) - x_j + f(s)) = u'(x^0 + f(s)) f'(s)$$

$$(\nabla_i C(\mathbf{x}) + f'(s)) \left(\sum_j p_j u'(C(\mathbf{x}) - x_j + f(s)) \right) = p_i u'(C(\mathbf{x}) - x_i + f(s)) + u'(x^0 + f(s)) f'(s)$$

solving for $\nabla_i C(\mathbf{x})$, we get

$$\begin{aligned} \nabla_i C(\mathbf{x}) &= \frac{p_i u'(C(\mathbf{x}) - x_i + f(s)) + u'(x^0 + f(s)) f'(s)}{\sum_j p_j u'(C(\mathbf{x}) - x_j + f(s))} - f'(s) \\ &= \frac{p_i u'(C(\mathbf{x}) - x_i + f(s)) + f'(s) \left(u'(x^0 + f(s)) - \sum_j p_j u'(C(\mathbf{x}) - x_j + f(s)) \right)}{\sum_j p_j u'(C(\mathbf{x}) - x_j + f(s))} \\ &= \frac{p_i u'(C(\mathbf{x}) - x_i + f(s))}{\sum_j p_j u'(C(\mathbf{x}) - x_j + f(s))} \end{aligned}$$

because by construction

$$u(x^0 + f(s)) = \sum_j p_j u(C(\mathbf{x}) - x_j + f(s))$$

and so their derivatives are also equal. Therefore

$$\sum_i \nabla_i C(\mathbf{x}) = \frac{\sum_i p_i u'(C(\mathbf{x}) - x_i + f(s))}{\sum_j p_j u'(C(\mathbf{x}) - x_j + f(s))} = 1$$

and so

$$\lim_{s \rightarrow \infty} \sum_i \mathcal{M}_i = \sum_i \nabla_i C(\mathbf{x}) + n f'(s) + n g'(s) = 1 + 0 + 0 = 1$$

■

5.2.2. Unbounded depth. In this section, we show that under additional, somewhat stronger, conditions a constant-utility profit-charging market maker also has unbounded depth. This result is similar to the one in Othman and Sandholm [2011a]. However, that market maker only showed unbounded depth when both u and f were logarithmic functions. Our result in this section is for a much broader class of utility, liquidity, and profit functions.

The first condition we require is that the utility function employed has *vanishing absolute risk aversion*. To our knowledge this concept was first defined in Caballé and Pomansky [1996].

Definition 9. A utility function has *vanishing absolute risk aversion* if it is twice-differentiable and

$$\lim_{c \rightarrow \infty} -\frac{u''(c)}{u'(c)} = 0$$

Most utility functions that are standard in the literature have this property. For instance, every utility function that has Constant Relative Risk Aversion (CRRA) [Mas-Colell et al. 1995] also has vanishing absolute risk aversion. An example of a utility function that does not have vanishing absolute risk aversion is $u(x) = -e^{-x}$.

As we have discussed, the *depth* of a market is the degree to which large bets move marginal prices. A very deep market will clear large orders without moving the marginal price. This leads us to the following definition.

Definition 10. Consider a market maker \mathcal{M} with a twice-differentiable price response, and let \mathcal{M}_{ii} denote the marginal change in the marginal price of a bet on the i -th event. We say that \mathcal{M} has *unlimited depth* if $\mathcal{M}_{ii} = 0$ for all i .

PROPOSITION 4. *Let f be a liquidity function with the property that $\lim_{s \rightarrow \infty} f(s) = \infty$. Then in any constant-utility profit-charging market maker, the arguments to the utility function grow arbitrarily large with s . Formally, for any \mathbf{x} and i*

$$\lim_{s \rightarrow \infty} C(\mathbf{x}) - x_i + f(s) = \infty$$

PROOF. Suppose not. Then there exists some finite bound B such that, for all i

$$\lim_{s \rightarrow \infty} C(\mathbf{x}) - x_i + f(s) < B$$

but then, because u is strictly increasing,

$$\lim_{s \rightarrow \infty} \sum_i p_i u(C(\mathbf{x}) - x_i + f(s)) < u(B)$$

But by construction,

$$\sum_i p_i u(C(\mathbf{x}) - x_i + f(s)) = u(x^0 + f(s))$$

and $\lim_{s \rightarrow \infty} x^0 + f(s) = \infty$. Therefore, there exists some S such that for all $s > S$,

$$x^0 + f(s) > B$$

At such s , we have

$$\sum_i p_i u(C(\mathbf{x}) - x_i + f(s)) < u(B) < u(x^0 + f(s))$$

but this is a contradiction because these values must be equal by construction. ■

PROPOSITION 5. *Let u be a utility function with vanishing absolute risk aversion, and let f be a liquidity function with the property that $\lim_{s \rightarrow \infty} f(s) = \infty$, and let the liquidity function f and profit function g have the property that*

$$\lim_{s \rightarrow \infty} f'(s) = \lim_{s \rightarrow \infty} f''(s) = \lim_{s \rightarrow \infty} g''(s) = 0$$

Then the constant-utility profit charging market maker formed by u , f , and g has unlimited depth as s gets large. Formally, $\lim_{s \rightarrow \infty} \mathcal{M}_{ii} = 0$.

PROOF. By definition $\mathcal{M}_{ii} = \nabla_{ii}^2 C(\mathbf{x}) + f''(s) + g''(s)$. Just as in the proof of Proposition 3, since

$$\lim_{s \rightarrow \infty} f''(s) = \lim_{s \rightarrow \infty} g''(s) = 0$$

the interesting term as s gets large is the cost function term.

For notational simplicity, define $s_i \equiv C(\mathbf{x}) - x_i + f(s)$. Then

$$\nabla_i C(\mathbf{x}) = \frac{p_i u'(s_i)}{\sum_j p_j u'(s_j)}$$

and taking the partial derivative with respect to i of both sides yields

$$\nabla_{ii}^2 C = \frac{p_i u''(s_i) s'_i \left(\sum_j p_j u'(s_j) \right)}{\left(\sum_j p_j u'(s_j) \right)^2} - \frac{p_i u'(s_i) \left(\sum_j p_j u''(s_j) s'_j \right)}{\left(\sum_j p_j u'(s_j) \right)^2} \quad (4)$$

We need to show that this approaches zero as s gets large.

First, $C(\mathbf{x})$ is a convex function because it is implicitly defined as an argument to equalize a concave utility function [Boyd and Vandenberghe 2004]. Therefore $\nabla_{ii}^2 C(\mathbf{x}) \geq 0$.

Now consider the s'_i terms. By construction $s'_{j \neq i} = \nabla_j C(\mathbf{x}) + f'(s)$ and $s'_i = \nabla_i C(\mathbf{x}) - 1 + f'(s)$. Therefore $1 \geq \lim_{s \rightarrow \infty} s'_{j \neq i} \geq 0 \geq \lim_{s \rightarrow \infty} s'_i \geq -1$. Thus we can bound the second derivative in the limit:

$$\lim_{s \rightarrow \infty} \nabla_{ii}^2 C(\mathbf{x}) \leq \lim_{s \rightarrow \infty} - \left[\frac{p_i u''(s_i) \left(\sum_j p_j u'(s_j) \right)}{\left(\sum_j p_j u'(s_j) \right)^2} + \frac{p_i u'(s_i) \left(\sum_j p_j u''(s_j) \right)}{\left(\sum_j p_j u'(s_j) \right)^2} \right]$$

where we have replaced all the s'_i in the first term of Equation 4 with -1 and all the s'_j in the second term of Equation 4 with 1, in order to make all of the positive terms as large as possible. (Recall that u is concave, and so $u'' \leq 0$.)

We will deal with the two terms of Equation 5.2.2 in succession, showing that the limit of each as s gets large is 0. We can simplify the first term immediately by canceling out the like term from the numerator and denominator, leaving

$$\lim_{s \rightarrow \infty} - \frac{p_i u''(s_i)}{\sum_j p_j u'(s_j)}$$

but then

$$\lim_{s \rightarrow \infty} - \frac{p_i u''(s_i)}{\sum_j p_j u'(s_j)} \leq \lim_{s \rightarrow \infty} - \frac{p_i u''(s_i)}{p_i u'(s_i)} = 0$$

because u has vanishing absolute risk aversion and because by Proposition 4, we have that $\lim_{s \rightarrow \infty} s_i = \infty$.

Now, we address the second term of Equation 5.2.2

$$\lim_{s \rightarrow \infty} - \frac{p_i u'(s_i) \left(\sum_j p_j u''(s_j) \right)}{\left(\sum_j p_j u'(s_j) \right)^2}$$

We will split this term into the product of two terms. First

$$\frac{p_i u'(s_i)}{\left(\sum_j p_j u'(s_j) \right)}$$

is just $\nabla_i C(\mathbf{x})$, and is therefore no larger than 1. Then because u has vanishing absolute risk aversion, and because as s gets large the s_j get large

$$\lim_{s \rightarrow \infty} - \frac{\left(\sum_j p_j u''(s_j)\right)}{\left(\sum_j p_j u'(s_j)\right)} = 0$$

Consequently,

$$\lim_{s \rightarrow \infty} - \frac{p_i u'(s_i) \left(\sum_j p_j u''(s_j)\right)}{\left(\sum_j p_j u'(s_j)\right)^2} = \lim_{s \rightarrow \infty} \frac{p_i u'(s_i)}{\left(\sum_j p_j u'(s_j)\right)} \cdot - \frac{\left(\sum_j p_j u''(s_j)\right)}{\left(\sum_j p_j u'(s_j)\right)} \leq 1 \cdot 0 = 0$$

Putting together all of our relations, we have

$$0 \leq \lim_{s \rightarrow \infty} \nabla_{ii}^2 C(\mathbf{x}) \leq 0$$

and so $\lim_{s \rightarrow \infty} \nabla_{ii}^2 C(\mathbf{x}) = 0$, and therefore

$$\lim_{s \rightarrow \infty} \mathcal{M}_{ii} = \lim_{s \rightarrow \infty} \nabla_{ii}^2 C(\mathbf{x}) + f''(s) + g''(s) = 0 + 0 + 0 = 0$$

■

5.3. Revenue bounds

To our knowledge, the only previous formal study of profit-charging behavior within the standard automated market making framework is Othman et al. [2010]. The automated market maker in that paper could enter a state of *outcome-independent profit*, which means that regardless of the realized outcome, the agent would book a profit. However, this condition only holds if the market terminates in that state. It is entirely possible, and in some settings virtually assured, that the market will leave the outcome-independent profit state and cause the market maker to run at a net loss.

To explore this notion further, recall that the market maker in Othman et al. [2010] is defined by the cost function

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

where $b(\mathbf{x}) = \alpha \sum_i x_i$. To see how outcome-independent profit works, consider a simple two-event market where $\alpha = 0.05$ and the market starts from $\mathbf{x}^0 = (1, 1)$. Imagine that the market maker takes two bets, the first a 50 dollar payout on the first event and the second a 50 dollar payout on the second event. The market maker is therefore in the state $(51, 51)$ and is in a state of outcome-independent profit, because no matter whether ω_1 or ω_2 is realized, the market maker books a profit of $C(51, 51) - 51 - C(\mathbf{x}^0) \approx 2.5 > 0$. However, imagine that another trader comes along and places an additional 50 dollar payout bet on the first event. Now the market maker is in state $(101, 51)$. If the market terminates and ω_1 is realized the market maker clears $C(101, 51) - 101 - C(\mathbf{x}^0) \approx -1.1 < 0$. So, the final bet has made the market maker exit the state of outcome-independent profit.

This problem is especially likely to arise in practice in settings where more information is revealed over time. For instance, consider a sports game in which betting is left open as the game proceeds. Once a clear winner emerges, traders will likely buy

many shares from the market maker in that team, causing the market maker to exit the state of unconditional profit.

There is an important loophole to this argument though. If the automated market maker is not designated with a formal, contractual role that necessitates its activity in a market, e.g., if it is being employed as the automated agent of a trader, then the market maker can circumvent this failure by simply ceasing to trade once it enters a state of outcome-independent profit, or alternately refusing to take any bet that would cause it to exit the state of outcome-independent profit.

In many settings, however, market makers are required to keep a presence in the market, e.g., if they are a designated market maker. For these settings, the notion of *unconditional profit* is more important. We say that a market maker enters a state of unconditional profit if it will make money regardless of the realized outcome *and* the future actions of traders. What is important to note about unconditional profit, as opposed to outcome-independent profit, is that once a market maker enters a state of unconditional profit it can never exit that state.

For example, consider a constant-utility profit-charging market maker using a barrier utility function. If the market maker's prior is positive, and if the profit function has collected more money than the initial value used to seed the liquidity function (i.e., if for all i , $p_i > 0$ and $g(s) > x^0$), then the market maker has entered the state of unconditional profit. This is because the market maker has no money pumps, loses no more than x^0 , and has collected a profit of more than this worst-case loss.

6. INSTANTIATING THE THEORY

In this section we collect the complete set of conditions we require on the various parts of the profit-charging automated market maker. We give several examples of functions that satisfy these conditions; any mix of these choices will produce an automated market maker with desirable qualities, differing only in the specifics of the initial amount of market depth, marginal bid/ask spread, and so on.

The prior distribution should be positive everywhere ($p_i > 0$) or the market maker will have unbounded worst-case loss. The uniform distribution is one option, but if the market maker has beliefs over the future these should be used, subject to never setting the probability of any event to zero. (If it is truly impossible for event ω_j to occur, then p_j can be set to zero while still maintaining bounded loss.)

The distance function d should be continuous. This is because a discrete distance function will never be able to achieve a vanishing bid/ask spread. One simple suggestion is to have d be (a positive multiple of) the distance function implied by any \mathcal{L}_p norm.

The utility function u should be a barrier utility function (in order to bound worst-case loss in a simple way), and should have vanishing absolute risk aversion (in order to have unlimited depth). One family of utility functions that satisfy both of these requirements is the standard CRRA utility functions

$$u(x) = \frac{x^{1-\gamma}}{1-\gamma}$$

indexed by $\gamma \geq 1$. (Here, log is the limit case for $\gamma = 1$.)

The liquidity function f should have $f(0) = 0$ and

$$\lim_{s \rightarrow \infty} f'(s) = \lim_{s \rightarrow \infty} f''(s) = 0$$

and also $\lim_{s \rightarrow \infty} f(s) = \infty$, to ensure a vanishing bid/ask spread and unlimited depth. One simple class of functions with both of these properties is $f(s) = \alpha s^{1/\beta}$ for $\beta > 1$

and $\alpha > 0$. Another class of functions that has these properties is $f(s) = \alpha \log(s + 1)$, again for $\alpha > 0$.

The profit function g should have $g(0) = 0$ and

$$\lim_{s \rightarrow \infty} g'(s) = \lim_{s \rightarrow \infty} g''(s) = 0$$

for unlimited depth and vanishing bid/ask spread. Since the conditions on the liquidity function f are a strict superset of the conditions on the profit function g , any liquidity function could also serve as a profit function. An example of a valid profit function that would not be a valid liquidity function is for the profit function to collect at most a certain amount more than the worst-case loss bound, e.g., setting

$$g(s) = \xi x^0 (1 - 1/(s + 1))$$

for $\xi > 1$ so that the market maker will collect at most ξ times x^0 . Once the market maker has collected more than x^0 , it is guaranteed to have entered an unconditional profit state.

To get a perspective on the operation of a representative market maker, consider Figures 1 and 2, which depict the prices charged by the market maker and the profit made by the market maker, respectively. In this specific example, the number of events is $n = 2$ and the market maker's probability on each event is $p_i = 1/2$. The distance function is implied by the \mathcal{L}_2 norm of the payout vector. The utility function is $u(x) = -1/x$, the initial wealth is $x^0 = 10$, the liquidity function is $f(s) = 100(\log(s + 10000) - \log(10000))$, and the profit function is $g(s) = 0.6(\sqrt{s + 100} - 10)$. These functions were chosen by experimenting with the classes of functions we suggested previously until we found a combination that made particularly attractive plots.

In both plots, the x axis is the quantity the market maker holds on both events before the interaction of the agents. That is, a value of 10^2 on the x axis indicates the market maker is at the payout vector $(100, 100)$ before the interaction of the trader.

Figure 1 shows the marginal price on the first event and the total price charged by the market maker for a one-dollar payout on the first event as the quantity in the market increases. Observe both that the cost of a fixed size bet approaches the marginal price and also that the marginal price approaches a bid/ask spread of 0 (which, because $p_1 = 0.5$ and the two events have equal payouts, corresponds to a marginal price of 0.5).

Figure 2 shows the profit made by the market maker as the amount transacted increases. For $s \gtrsim 600$, the market maker enters a state of unconditional profit, because $g(s) > x^0 = 10$. This state is reached at a payout of at most about 430 on each event. If there is additional churn in the market—traders that sell their bets back to the market maker—the state of unconditional profit can be reached at smaller payout vectors. At an extreme, a state of unconditional profit can be made by an agent buying and selling back a single bet to the market maker a large number of times, each time paying a small amount into the profit function.

7. CONCLUSIONS

We constructed a new framework for automated market making agents that builds on the prior literature. Our market making agents are the first to simultaneously satisfy four key desiderata that have appeared in the literature: (1) bounded worst-case loss, (2) the ability to make a profit, (3) vanishing bid/ask spread, and (4) unlimited depth as trading volume grows large. All triples of these four properties have been achieved previously, but satisfying all four is not straightforward because several of the properties are oppositional.

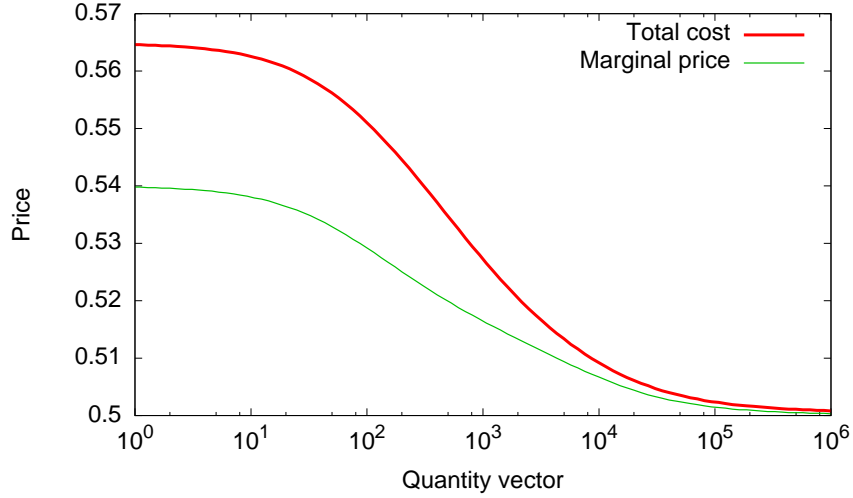


Fig. 1. Total cost for making a one-ticket bet and marginal prices for the market maker and setting given in the text. Observe that as the quantity transacted increases the bid/ask spread goes to zero (i.e., both lines go to 0.5). The x axis is log-scaled.

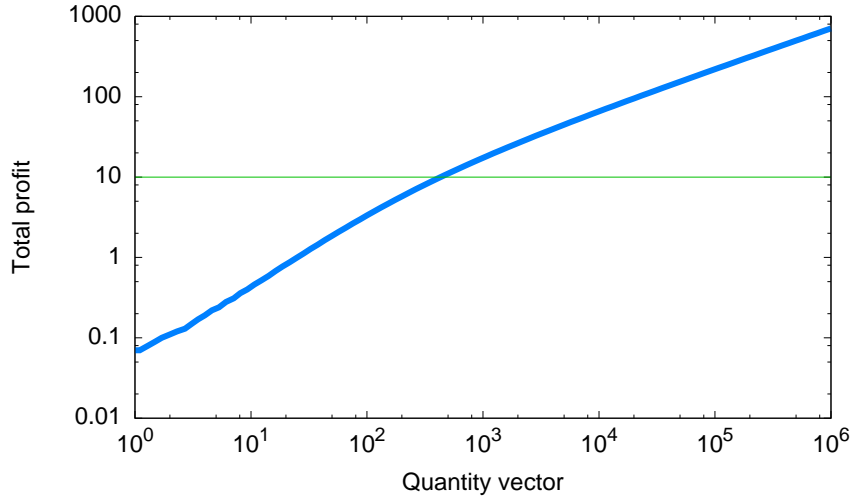


Fig. 2. The profit collected increases with the quantity transacted in the market. The straight line at 10 represents when the market maker enters a state of unconditional profit. Both axes are log-scaled.

The approach that enables us to satisfy all four desiderata is based on new framework of thinking about market makers. It extends constant-utility cost functions with two separate functions that are added to the prices quoted to the trader. The *liquidity function* uses its proceeds to increase the amount of liquidity provided by the market maker. The *profit function* represents a “lockbox” of savings that is separate from the rest of the market maker’s decision-making process.

Unlike most of the prior literature, interactions with our agents are path dependent. For most of the market makers in the literature, (including the *de facto* market maker in practice, the LMSR), a trader can buy and then immediately sell a contract from

the market maker at no cost. However, as we showed in Table I, real markets do not generally function this way: a trader that buys and then immediately sells a contract from the market maker will book a loss. Our market makers match real world market makers in their path dependence.

8. FUTURE RESEARCH

There are several extensions worth considering. First, we have heard the argument that it could make more sense to have smaller bid/ask spreads at the market's initiation, to encourage participation, and widen the bid/ask spread as the market matures. This scheme has an appealing intuition: essentially, early participants are being compensated for providing liquidity at the more fragile beginning stages of the market. One way to produce this effect with our market makers is to start s at a large value and decrement it over time (e.g., $s \leftarrow s - d(x, y)$ when the value is updated). This essentially runs the state variable s in reverse. However, as discussed in Section 2, this is counter to the way real markets function, and so adopting this methodology may be unintuitive for participants. There may be better ways of doing this, and we believe that our framework of profit and liquidity functions can facilitate such study.

Another extension is to extend our agents to handle not only market orders, but also limit orders—orders that execute only if certain price conditions are reached. Here, a possible direction seems to be to extend the convex optimization framework of Agrawal et al. [2009]. That work provided a simple, computationally efficient way to turn a cost function that is capable of only handling market orders into a cost function that can handle both market and limit orders. While adding a profit function to the Agrawal framework is trivial, incorporating a liquidity function appears to be more challenging, because the liquidity function is used in calculating prices. To be explicit, the trading volume affects the liquidity function, which affects the cost function, which affects the number of cleared orders, which affects the trading volume. Resolving this circularity seems like the most direct solution for our market making agents to handle persistent limit orders.

An extension of a different flavor is to consider the values we captured from real markets in Table I. To our knowledge, that table represents the first comparison of the prevalent liquidity in prediction markets versus equity markets. That table showed, somewhat surprisingly, that active prediction market contracts have a bid/ask spread roughly comparable to small cap equities, despite the fact that publicly traded equities have several orders of magnitude greater daily volume. This suggests that popular prediction markets may be more robust than would be expected from their small transaction volumes. It would be interesting to study whether this phenomenon holds more broadly, and to understand why.

REFERENCES

- J. Abernethy, Y. Chen, and J. W. Vaughan. An optimization-based framework for automated market-making. In *ACM Conference on Electronic Commerce (EC)*, 2011.
- S. Agrawal, E. Delage, M. Peters, Z. Wang, and Y. Ye. A unified framework for dynamic pari-mutuel information market design. In *ACM Conference on Electronic Commerce (EC)*, pages 255–264, 2009.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- J. Caballé and A. Pomansky. Mixed risk aversion. *Journal of Economic Theory*, 71(2): 485–513, 1996.
- M. Chakraborty, S. Das, A. Lavoie, M. Magdon-Ismail, and Y. Naamad. Instructor rating markets. In *Conference on Auctions, Market Mechanisms and Their Applications*

- (AMMA), 2011.
- Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 49–56, 2007.
- Y. Chen and D. M. Pennock. Designing markets for prediction. *AI Magazine*, 31(4), 2010.
- Y. Chen and J. W. Vaughan. A new understanding of prediction markets via no-regret learning. In *ACM Conference on Electronic Commerce (EC)*, pages 189–198, 2010.
- H. Demsetz. The cost of transacting. *The Quarterly Journal of Economics*, 82(1):33–53, 1968.
- S. Goel, D. Pennock, D. Reeves, and C. Yu. Yoopick: a combinatorial sports prediction market. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1880–1881, 2008.
- R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
- D. MacKenzie. *An engine, not a camera: How financial models shape markets*. The MIT Press, 2006.
- A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, New York, NY, 1995.
- M. O’Hara. *Market Microstructure Theory*. Blackwell Publishing, 1995.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. In *ACM Conference on Electronic Commerce (EC)*, pages 253–254, 2009.
- A. Othman and T. Sandholm. Automated market-making in the large: the Gates Hillman prediction market. In *ACM Conference on Electronic Commerce (EC)*, pages 367–376, 2010a.
- A. Othman and T. Sandholm. When Do Markets with Simple Agents Fail? In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 865–872, Toronto, Canada, 2010b.
- A. Othman and T. Sandholm. Automated market makers that enable new settings: Extending constant-utility cost functions. In *Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, 2011a.
- A. Othman and T. Sandholm. Liquidity-sensitive market makers via homogeneous risk measures. In *International Workshop On Internet And Network Economics (WINE)*, 2011b.
- A. Othman and T. Sandholm. Inventory-based versus prior-based options trading agents. *Algorithmic Finance*, to appear, 2012.
- A. Othman, D. M. Pennock, D. M. Reeves, and T. Sandholm. A practical liquidity-sensitive automated market maker. In *ACM Conference on Electronic Commerce (EC)*, pages 377–386, 2010.
- D. Pennock and R. Sami. Computational Aspects of Prediction Markets. In *Algorithmic Game Theory*, chapter 26, pages 651–674. Cambridge University Press, 2007.
- M. Peters, A. M.-C. So, and Y. Ye. Pari-mutuel markets: Mechanisms and performance. In *International Workshop On Internet And Network Economics (WINE)*, pages 82–95, 2007.
- N. Taleb. *The black swan: The impact of the highly improbable*. Random House Inc, 2007.
- U.S. Commodity and Futures Trading Commission and U.S. Securities & Exchange Commission. Findings regarding the market events of May 6, 2010. Technical report, Joint Advisory Committee on Emerging Regulatory Issues, September 2010.