

Market Making and Mean Reversion

Tanmoy Chakraborty
University of Pennsylvania
tanmoy@seas.upenn.edu

Michael Kearns
University of Pennsylvania
mkearns@cis.upenn.edu

ABSTRACT

Market making refers broadly to trading strategies that seek to profit by providing liquidity to other traders, while avoiding accumulating a large net position in a stock. In this paper, we study the profitability of market making strategies in a variety of time series models for the evolution of a stock's price. We first provide a precise theoretical characterization of the profitability of a simple and natural market making algorithm in the absence of any stochastic assumptions on price evolution. This characterization exhibits a trade-off between the positive effect of local price fluctuations and the negative effect of net price change. We then use this general characterization to prove that market making is generally profitable on *mean reverting* time series — time series with a tendency to revert to a long-term average. Mean reversion has been empirically observed in many markets, especially foreign exchange and commodities. We show that the slightest mean reversion yields positive expected profit, and also obtain stronger profit guarantees for a canonical stochastic mean reverting process, known as the *Ornstein-Uhlenbeck (OU) process*, as well as other stochastic mean reverting series studied in the finance literature. We also show that market making remains profitable in expectation for the OU process even if some realistic restrictions on trading frequency are placed on the market maker.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Theory, Economics, Algorithms

Keywords

Computational Finance, Algorithmic Trading, Market Making, Mean Reversion

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'11, June 5–9, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0261-6/11/06 ...\$10.00.

1. INTRODUCTION

A market maker is a firm, individual or trading strategy that always or often quotes both a buy and a sell price for a financial instrument or commodity, hoping to make a profit by exploiting the difference between the two prices, known as the *spread*. Intuitively, a market maker wishes to buy and sell equal volumes of the instrument (or commodity), and thus rarely or never accumulate a large net position, and profit from the difference between the selling and buying prices.

Historically, the chief purpose of market makers has been to provide liquidity to the market — the financial instrument can always be bought from, or sold to, the market maker at the quoted prices. Market makers are common in foreign exchange trading, where most trading firms offer both buying and selling rates for a currency. They also play a major role in stock exchanges, and historically exchanges have often appointed trading firms to act as official market makers for specific equities. NYSE designates a single market maker for each stock, known as the *specialist* for that stock. In contrast, NASDAQ allows several market makers for each stock. More recently, fast electronic trading systems have led trading firms to behave like market makers without formally being designated so. In other words, many trading firms attempt to buy and sell a stock simultaneously, and profit from the difference between buying and selling prices. We shall refer to such trading algorithms generally as *market making algorithms*.

In this paper, we analyze the profitability of market making algorithms. Market making has existed as a trading practice for a long time, and it has also inspired significant amount of empirical as well as theoretical research [9, 5, 10, 1, 2, 3]. Most of the theoretical models [5, 9, 2, 3] view market makers as dealers who single-handedly *create* the market by offering buying and selling prices, and there is no trading in their absence (that is, all trades must have the market maker as one of the parties). On the other hand, much of the empirical work has focused on analyzing the behavior of specialist market makers in NYSE, using historical trading data from NYSE [10, 1]. In contrast, our theoretical and empirical work studies the behavior of market making algorithms in both very general and certain specific price time series models, where trading occurs at varying prices even in the absence of the market maker. This view seems appropriate in modern electronic markets, where any trading party whatsoever is free to quote on both sides of the market, and officially designated market makers and specialists are of diminishing importance.

Market Making vs. Statistical Arbitrage.

Before describing our models and results, we first offer some clarifying comments on the technical and historical differences between market making and statistical arbitrage, the latter referring to the activity of using computation-intensive quantitative modeling to design profitable automated trading strategies. Such clarification is especially called for in light of the blurred distinction between traditional market-makers and other kinds of trading activity that electronic markets have made possible, and the fact that many quantitative hedge funds that engage in statistical arbitrage may indeed have strategies that have market making behaviors.

Perhaps the hallmark of market making is the willingness to always quote competitive buy and sell prices, but with the goal of *minimizing directional risk*. By this we mean that the market maker is averse to acquiring a large net long or short position in a stock, since in doing so there is the risk of large losses should the price move in the wrong direction. Thus if a market maker begins to acquire a large net long position, it would continue to quote a buy price, but perhaps a somewhat lower one which is less likely to get executed. Alternatively (or in addition), the strategy might choose to lower its sell quote in order to increase the chances of acquiring short trades to offset its net long inventory. In this sense, pure market making strategies have no “view” or “opinion” on which direction the price “should” move — indeed, as we shall show, the most profitable scenario for a market maker is one in which there is virtually no overall directional movement of the stock, but rather a large amount of non-directional volatility.

In contrast, many statistical arbitrage strategies are the opposite of market making in that they deliberately want to make directional bets — that is, they want to acquire large net positions because they have a prediction or model of future price movement. To give one classic example, in the simplest form of pairs trading, one follows the prices of two presumably related stocks, such as Microsoft and Apple. After normalizing the prices by their historical means and variances, one waits for there to be a significant gap in their current prices — for instance, Apple shares becoming quite expensive relative to Microsoft shares given the historical normalization. At this point, one takes a large short position in Apple and an offsetting large long position in Microsoft. This amounts to a bet that the prices of the two stocks will eventually return to their historical relationship: if Apple’s share price falls and Microsoft’s rises, both positions pay off. If the gap continues to grow, the strategy incurs a loss. If both rise or both fall without changing the gap between, there is neither gain nor loss. The important point here is that, in contrast to market making, the source of profitability (or loss) are directional bets rather than price volatility.

Theoretical Model and Results.

We first summarize our theoretical models and our three main theoretical results. We assume there is an exogenous market where a stock can be bought and sold at prices dictated by a given time series process. At any given point of time, there is a single *exogenous asset price* at which the stock can both be bought as well as sold. The price evolution in the market is captured by the time series. The market making algorithm is an online decision process that

can place buy and sell *limit orders* with some quoted limit order prices at any time, and may also cancel these orders at any future time. For simplicity, we assume that each order requests only one share of the stock (a trader may place multiple orders at the same price). If at any time after placing the order and before its cancellation, the asset price of the stock equals or exceeds (respectively, falls below) the quoted price on a sell order (respectively, buy order), then the order gets *executed at the quoted price*, *i.e.* the trader pays (respectively, gains) one share and gains (respectively, pays) money equal to the price quoted on the order. We shall refer to the net volume of the stock held by a trader at a given point of time as *inventory*. Note that inventory may be positive (net long position) or negative (net short position). To evaluate the profit made by a market making algorithm, we shall fix a time horizon when the algorithm must liquidate its inventory at the current asset price.

Our first and most general theoretical result is a succinct and *exact* characterization of the profit obtained by a simple market-making algorithm, given *any* asset price time series, in the model above. If the sum of absolute values of all *local price movements* (defined below) is K , and the difference between opening and closing prices is z , we show that the profit obtained is exactly $(K - z^2)/2$. The positive term K can be viewed as the product of the average volatility of the price and the duration for which the algorithm is run. The negative term z^2 captures the *net change in price* during the entire trading period. Thus this characterization indicates that *market making is profitable when there is a large amount of local price movement, but only a small net change in the price*. This observation matches a common intuition among market makers, and provides a theoretical foundation for such a belief. An unbiased random walk (or Brownian motion) provides a *boundary* of profitability — the algorithm makes zero expected profit (as do all trading algorithms), while any stochastic price process whose closing price has comparatively less variance from the opening price makes positive expected profit. The last observation leads to our second result.

Mean Reversion.

We next exhibit the benefit of obtaining a succinct and exact expression for profit by applying it to some classes of stochastic time series that help in understanding the circumstances under which the algorithm is profitable. We identify a natural class of time series called *mean-reverting* processes whose properties make our market making algorithm profitable in expectation. A stochastic price series is considered to be *reverting* towards its *long-term mean* μ if the price shows a downward trend when greater than μ and upward trend when less than μ . Prices of commodities such as oil [11, 13] and foreign exchange rates [8] have been empirically observed to exhibit mean reversion. Mean-reverting stochastic processes are studied as a major class of price models, as a contrast to stochastic processes with directional drift, or with no drift, such as Brownian motion. One widely studied mean-reverting stochastic process is the *Ornstein-Uhlenbeck process* [7].

Formally, our second result states that our market making algorithm has expected positive profit on any random walk that reverts towards its opening price. This result is quite revealing — it holds if the random walk shows even the slightest mean reversion, regardless of how complex the

process may be (for instance, its evolution may depend not only on the current price, but also on the historical prices in an arbitrary way, as well as the current time). It identifies mean reversion as the natural property which renders market making profitable.

Our third result shows that simple market making algorithms yield stronger profit guarantees for specific mean-reverting processes. As an example, we consider the Ornstein-Uhlenbeck (OU) process. If the price series follows this process, we show that a simple market making algorithm is profitable when run for a sufficiently long time. Moreover, the profit grows *linearly* with the duration for which the algorithm is run, and the profit guarantees hold not only in expectation, but with high probability. We prove this by showing that while $\mathbf{E}[K]$ grows linearly with time, $\mathbf{E}[z^2]$ is upper bounded by a constant. Unlike our second result, we do not need the assumption that the price series begins at the long-term mean — the initial price appears in the upper bound on $\mathbf{E}[z^2]$.

We also show an analogous result for another mean reverting process that has been studied in the finance literature, a model studied by Schwartz [14]. In this model, the local volatility is a linear function of price, while the OU process models volatility as a constant.

We remark that the results outlined above assume a model where the market maker can place and cancel orders as frequently as it likes, and in fact our algorithm does so after every change in the exogenous asset price. In practice, however, a market maker cannot react to every change, since the asset price may change with every trade in the market (which may or may not involve this particular market maker), and the market maker may not be able to place new limit orders after every trade of a rapidly traded stock. So we also analyze the profitability of our market making algorithm when it is allowed to change its orders only after every L steps, by simulating our algorithm on random samples from the OU process. If the price series is the OU process, we show that the expected profit continues to grow linearly with time.

Other Related Work.

To our knowledge, no previous work has studied market making in an exogenously specified price time series model. Most of the theoretical work, as mentioned before, considers a single dealer model where all trades occurred through the market maker at its quoted prices [5, 9, 2, 3]. This includes the well-known Glosten-Milgrom model for market making [9]. On the other hand, there has been a fair amount of work in algorithmic trading, especially statistical arbitrage, that assumes an exogenous price time series. The closest line of research to our work in this literature is the analysis of pair trading strategies under the assumption that the price difference between the two equities show mean reversion (e.g. [4, 12]). As discussed before, such strategies are qualitatively very different from market making strategies. Moreover, most algorithmic trading work, to our knowledge, either analyze price series given by very specific stochastic differential equations (similar to Sections 3.1 and 3.2 of our paper), or empirically analyze these algorithms against historical trading data (e.g. [6]). In contrast, we also give profit guarantees for the weakest of mean reversion processes without assuming a specific form (Theorem 3.1), and in fact derive an exact expression for arbitrary price series (Theorem

2.1), inspired from the notion of worst-case analysis in theoretical computer science.

2. A GENERAL CHARACTERIZATION

We first describe our theoretical model formally. We assume that all events occur at discrete time steps $0, 1, 2 \dots T$, where T is the time horizon when the market making algorithm must terminate. There is an asset price P_t of the stock at every time step $0 \leq t \leq T$. Thus $P_0, P_1 \dots P_T$ is the *asset price time series*. We assume that all prices are *integral* multiples of a basic unit of money (regulations in NYSE/NASDAQ currently require prices to be integral multiples of a penny).

A trading algorithm may place and cancel (possibly both) buy and sell orders at any of these time steps, and each order requests a single share at a quoted limit order price Y . A buy (respectively, sell) order at price Y placed at time t gets *executed* at the earliest time $t' > t$ such that $P_{t'} \leq Y$ (respectively, $P_{t'} \geq Y$), provided that the order is not canceled before t' . Any buy order placed by our algorithms will quote a price $Y < P_t$, and so cannot get executed at time t itself; the same applies to sell orders as well. If the buy (respectively, sell) order gets executed (at some time in the future), then the algorithm pays (respectively, earns) Y units of money due to the execution. We assume *no market impact* — the orders placed by the algorithm (which may get executed) do not affect the prices in the future. We leave to future work the important topic of incorporation of market impact into our results.

The *inventory* held by the algorithm at time t is the number of shares held by it at that time. The inventory is incremented upon every executed buy order, and decremented upon every executed sell order. The initial inventory is, naturally, assumed to be zero. At time T , the algorithm must liquidate its accumulated inventory at the current asset price (alternatively, one may view it as the evaluation of the portfolio at time T). If the algorithm has inventory x (x may be negative) at time T , then it earns or pays xP_T from the liquidation. The *profit* obtained by the algorithm is then its net cash position (positive or negative) after liquidation.

A trading algorithm is considered to be an online process — it makes its decisions (about placing and canceling orders) at time t after observing the price series up to (and including) time t . The algorithm may or may not have additional information about the price series. Note that we assume *no latency* and *arbitrary frequency* (we relax the frequency assumption in our simulations) — the algorithm can look at current prices and place buy and sell orders instantaneously, and it can do so as frequently as it wishes. Again, we leave the relaxation of these unrealistic assumptions for future work.

Market Making Algorithms.

The basic class of market making algorithms that we consider is the following: At time t , the algorithm cancels all unexecuted orders, and places new buy orders at prices $Y_t, Y_t - 1, Y_t - 2 \dots Y_t - C_t$ and new sell orders at prices $X_t, X_t + 1, X_t + 2 \dots X_t + C_t$, where $Y_t < X_t$ and C_t is a non-negative integer. Such *ladders* of prices are set up to ensure that large sudden fluctuations in price causes a proportionally large volume of executions. C_t is called the *depth* of the price ladder at time t , and intuitively, the algorithm believes that




Figure 1: Matched and unmatched trades

the price fluctuation $|P_{t+1} - P_t|$ shall not exceed C_t . Un-executed orders are canceled and fresh orders are placed at every time step (changes are necessary only if $X_{t+1} \neq X_t$ or $Y_{t+1} \neq Y_t$). Thus the algorithm is determined by the choices of X_t , Y_t and C_t for all t , and these choices may be made after observing the price movements up to time t .

We begin by presenting our basic result for a simple market making algorithm, that sets $X_t = P_t + 1$ and $Y_t = P_t - 1$.

THEOREM 2.1. *Let $P_0, P_1 \dots P_T$ be an asset price time series. Let $K = \sum_{t=1}^T |P_t - P_{t-1}|$, and let $z = P_T - P_0$. Suppose that $|P_{t+1} - P_t| \leq D_t \forall t$, where D_t is known to the algorithm at time t . If the market making algorithm, that sets $X_t = P_t + 1$, $Y_t = P_t - 1$ and ladder depth $C_t = D_t$, is run on this price series, then the inventory liquidated at time T is $-z$, and the profit is $(K - z^2)/2$.*

PROOF. Note that at any time step $t > 0$, at least one order gets executed if the price changes. Moreover, the number of orders executed at time $t \geq 1$ is $P_{t-1} - P_t$ (a negative value indicates that shares were sold). The statement holds as long as $|P_{t-1} - P_t| \leq C_{t-1}$, which is true by assumption. Thus K is equal to the total number of orders that gets executed. Moreover, the size of inventory held by the algorithm at time t is $P_0 - P_t$. We shall construct disjoint pairs of all but $|z|$ of the executed orders, such that each pair of executions comprises an executed buy and a sell execution, and the price of the buy order is 1 less than the price of the sell order, so that each such pair can be viewed as giving a profit of 1.

For $p > P_0$, we pair each sell order, priced at p , that gets executed when price increases to p or more, with the executed buy order, priced at $p - 1$, that gets executed at the earliest time in the future when the price falls back to $p - 1$ or less (if the price ever falls to $p - 1$ again). Note that these pairs are disjoint, since between every rise of the price from $p - 1$ to p , the price must obviously fall to $p - 1$. Similarly, for $p < P_0$, we pair each executed buy order when the price decreases to p or less, with the executed sell order at the earliest time in the future when the price increases to $p + 1$ or more (if it exists).

We claim that only $|z|$ executions remain unmatched (see Figure 2). If $z > 0$, then the only executions that remain unmatched are the sell orders executed when the price increases to p and never again falls below p : for each $P_0 + z \geq p > P_0$, there is one such executed order. During liquidation at time

T , these unmatched sell orders are matched by buying z shares at price $P_T = P_0 + z$. The total loss during liquidation is $\sum_{p=P_0+1}^{P_0+z} ((P_0 + z) - p) = z(z - 1)/2$. Since there are $K - z$ paired executions, the profit obtained from them is $(K - z)/2$. Hence the net profit is $(K - z - z(z - 1))/2 = (K - z^2)/2$. A symmetric argument holds for $z < 0$. \square

Note that it is typically reasonable to assume that $|P_{t+1} - P_t| < C$ for some large enough C , since it is unlikely that the price of a stock would change by more than a few dollars within a few seconds. In that case, one may set $C_t = C \forall t$.

3. MEAN REVERSION MODELS

In this section, we use Theorem 2.1 to relate profitability to mean reversion.

DEFINITION 3.1. *A unit-step walk is a series $P_0, P_1 \dots P_T$ such that $|P_{t+1} - P_t| \leq 1 \forall T > t \geq 0$. A stochastic price series $P_0, P_1 \dots P_T$ is called a random walk if it is a unit-step walk almost surely. We say that a random walk is unbiased if $\Pr [P_{t+1} - P_t = 1 | P_t, P_{t-1} \dots P_0] = 1/2$, for all unit-step walks $P_t \dots P_0$, for all $T > t \geq 0$.*

We say a random walk is mean-reverting towards μ if

$$\Pr [P_{t+1} - P_t = 1 | P_t = x, P_{t-1} \dots P_0] \geq \Pr [P_{t+1} - P_t = -1 | P_t = x, P_{t-1} \dots P_0]$$

for all $x \leq \mu$, and

$$\Pr [P_{t+1} - P_t = 1 | P_t = y, P_{t-1} \dots P_0] \leq \Pr [P_{t+1} - P_t = -1 | P_t = y, P_{t-1} \dots P_0]$$

for all $y \geq \mu$, for all $t, P_{t-1}, P_{t-2} \dots P_0$ such that $P_0, \dots P_t$ is a unit-step walk, and at least one of these inequalities for $t < T$ is strict (i.e., it is not an unbiased random walk).

Note that all trading algorithms yield zero expected profit on an unbiased random walk. This is because the profit F_t of the algorithm, if its inventory were liquidated at time t , is a martingale, irrespective of the number of shares bought or sold at each time step, and so the expected profit is $\mathbf{E}[F_T] = \mathbf{E}[F_0] = 0$.

THEOREM 3.1. *For any random walk $P_0, P_1 \dots P_T$ that is mean-reverting towards $\mu = P_0$, the expected profit of the market making algorithm that sets $X_t = P_t + 1$ and $Y_t = P_t - 1$ (any $C_t \geq 0$ suffices) is positive.*

PROOF. Since the price does not change by more than 1 in a time step, the market making algorithm need not set a ladder of prices. By Theorem 2.1, the expected profit is $\mathbf{E}[(K - z^2)/2]$. Let $K_t = \sum_{i=1}^t |P_i - P_{i-1}|$, and let $z_t = P_t - P_0$. We show by induction on t that $\mathbf{E}[K_t] \geq \mathbf{E}[z_t^2]$ for all t . For $t = T$, this would imply positive expected profit for our algorithm. Without loss of generality, we assume that $P_0 = \mu = 0$.

For $t = 0$, the statement is trivially true, since $K_t = z_t = 0$. Suppose it is true for some $t \geq 0$, then we can show that it is true for $t + 1$. Let \mathcal{F}_t denote the set of all unit-step walks such that $P_0 = \mu = 0$. For $s \in \mathcal{F}_t$, let $\alpha(s) = \Pr [P_{t+1} - P_t = 1 | P_t, P_{t-1} \dots P_0]$, and let $\beta(s) = \Pr [P_{t+1} - P_t = -1 | P_t, P_{t-1} \dots P_0]$. Also, let $\Pr [s]$ denote the probability that the first t steps of this random walk is s . Then we have

$$\begin{aligned} \mathbf{E}[K_{t+1}] &= \mathbf{E}[K_t + |P_{t+1} - P_t|] \\ &= \mathbf{E}[K_t] + \sum_{s \in \mathcal{F}_t} \Pr [s] (\alpha(s) + \beta(s)) \end{aligned} \quad (1)$$

$$\begin{aligned}
\mathbf{E}[z_{t+1}^2] &= \mathbf{E}[P_{t+1}^2] \\
&= \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] (\alpha(s)(P_t + 1)^2 + \beta(s)(P_t - 1)^2 \\
&\quad + (1 - \alpha(s) - \beta(s))P_t^2) \\
&= \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] (P_t^2 + \alpha(s) + \beta(s) + 2P_t(\alpha(s) - \beta(s))) \\
&= \mathbf{E}[P_t^2] + \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] (\alpha(s) + \beta(s)) \\
&\quad + 2 \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] P_t(\alpha(s) - \beta(s)) \\
&\leq \mathbf{E}[K_t] + \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] (\alpha(s) + \beta(s)) \\
&\quad + 2 \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] P_t(\alpha(s) - \beta(s)) \\
&\quad \text{(by induction hypothesis)} \\
&= \mathbf{E}[K_{t+1}] + 2 \sum_{s \in \mathcal{F}_t} \mathbf{Pr}[s] P_t(\alpha(s) - \beta(s)) \\
&\quad \text{(by Equation 1)}
\end{aligned}$$

It suffices to show that that $P_t(\alpha(s) - \beta(s)) \leq 0$ for all $s \in \mathcal{F}_t$. This follows immediately from the definition of a mean-reverting random walk: if $P_t > P_0 = 0$, then $\alpha(s) < \beta(s)$, and if $P_t < P_0 = 0$, then $\alpha(s) > \beta(s)$. Thus we have proved the induction hypothesis for $t + 1$.

Finally, for the smallest t such that for some $s \in \mathcal{F}_t$ we have $\alpha(s) \neq \beta(s)$, the inequality in the induction hypothesis becomes strict at $t + 1$, *i.e.* $\mathbf{E}[K_{t+1}] > \mathbf{E}[z_{t+1}^2]$, and so the expected profit for $T > t$ is strictly positive. \square

3.1 Ornstein-Uhlenbeck Processes

One well-studied mean-reverting process is a continuous time, real-valued stochastic process known as the *Ornstein-Uhlenbeck (OU) process* [7]. We denote this process by Q_t . It is usually expressed by the following stochastic differential equation:

$$dQ_t = -\gamma(Q_t - \mu)dt + \sigma dW_t,$$

where W_t is a standard Brownian motion, and γ, σ are positive constants, and $\gamma < 1$. The value μ is a constant around which the price fluctuates — it is called the *long term mean* of the process. The coefficient of dt is called *drift*, while that of dW_t is called *volatility*. Observe that the drift is negative for $P_t > \mu$ and positive for $P_t < \mu$ — this is why the process tends to revert towards μ whenever it is far from it. γ is the rate of mean reversion. The OU process is memoryless (distribution of Q_t given Q_0 is the same as distribution of Q_{t+x} given Q_x), and given an opening value Q_0 , the variable Q_t is known to be normally distributed, such that

$$\begin{aligned}
\mathbf{E}[Q_t] &= \mu + (Q_0 - \mu)e^{-\gamma t}, \\
\mathbf{Var}[Q_t] &= \frac{\sigma^2}{2\gamma}(1 - e^{-2\gamma t})
\end{aligned} \tag{2}$$

Now we consider the OU process Q_t as a price series in the unique model, and analyze profitability of our algorithm. However, since the OU process is a continuous time real-valued process, we need to define a natural restriction to a discrete integral time series that conforms to our theoretical model. We achieve this by letting P_t to be the nearest

integer to Q_t , for all non-negative integers t . The rounding is practical since in reality prices are not allowed to be real-valued, and further, our algorithm reacts to only integral changes in price. We shall analyze our algorithm on P_t .

A significant hindrance in applying Theorem 2.1 to the OU process is that the jumps $|P_{t+1} - P_t|$ are not necessarily bounded by some constant C , so we have to put some effort into determining C_t . Since the OU process is memoryless, Equation 2 implies that given Q_t , Q_{t+1} is normally distributed with expectation $\mu + (Q_t - \mu)e^{-\gamma}$ and variance less than $\sigma^2/2\gamma$, so if we set $C \gg \frac{\sigma^2}{2\gamma}\sqrt{\ln T}$ and then set $C_t = \mathbf{E}[|Q_{t+1} - Q_t| | Q_t] + C$, then the probability that the price jump at any time exceeds the depth of the ladder is vanishing, and such events do not contribute significantly to the expected profit if we simply stop the algorithm when such an event occurs.

THEOREM 3.2. *Let $P_0, P_1 \dots P_T$ be a price series obtained from an OU process $\{Q_t\}$ with long-term mean μ . Then the market making algorithm that sets $X_t = P_t + 1$, $Y_t = P_t - 1$ and $C_t = \mathbf{E}[|Q_{t+1} - Q_t| | Q_t] + 10\frac{\sigma^2}{2\gamma}\sqrt{\ln T}$ yields long term expected profit $\Omega(\sigma T - \sigma^2/2\gamma - (\mu - Q_0)^2)$.*

PROOF. It is easy to show, as outlined above, that the contribution of events where a price jump larger than C_t occurs to the expected profit is negligible. We restrict our attention to the event when no such large jump occurs. By Theorem 2.1, the profit on a sample series is $(K - z^2)/2$, where $K = \sum_{t=1}^T |P_t - P_{t-1}|$, and $z = P_T - P_0$. The result follows by giving a lower bound on $\mathbf{E}[K]$ and an upper bound on $\mathbf{E}[z^2]$.

Let us derive a lower bound on $\mathbf{E}[|Q_{t+1} - Q_t|]$. Note that this quantity is equal to $\mathbf{E}[|Q'_1 - Q'_0| | Q'_0 = Q_t]$, where Q'_i is an identical but independent OU process. This is because the OU process is Markov, and future prices depend only on the current price. Since $\gamma < 1$, so Equation 2 implies that given Q_t , Q_{t+1} is normally distributed with variance greater than $\sigma^2/4$, since $\frac{1-e^{-2\gamma}}{2\gamma} > 1/4$ when $\gamma < 1$. Hence $\mathbf{E}[|Q_{t+1} - Q_t|]$ is at least $\sigma/4$ (using properties of a folded normal distribution). Since P_t is obtained by rounding Q_t , we have $|P_{t+1} - P_t| > |Q_{t+1} - Q_t| - 2$. Thus for large enough σ (see comments at the end of the theorem), we get that $\mathbf{E}[K] = \Omega(\sigma T)$.

$\mathbf{E}[z^2]$ is approximated well enough by $\mathbf{E}[(Q_T - Q_0)^2]$, since $|z - (Q_T - Q_0)| < 2$ for all possible realizations. Again, Equation 2 implies that $Q_T - Q_0$ has mean $\mu + (Q_0 - \mu)e^{-\gamma T} - Q_0 = (\mu - Q_0)(1 - e^{-\gamma T})$ and variance $\sigma^2(1 - e^{-2\gamma T})/2\gamma$. Thus, we have

$$\begin{aligned}
\mathbf{E}[(Q_T - Q_0)^2] &= \mathbf{Var}[Q_T - Q_0] + \mathbf{E}[Q_T - Q_0]^2 \\
&= \frac{\sigma^2(1 - e^{-2\gamma T})}{2\gamma} \\
&\quad + (\mu - Q_0)^2(1 - e^{-\gamma T})^2 \\
&< \frac{\sigma^2}{2\gamma} + (\mu - Q_0)^2
\end{aligned}$$

Thus, $\mathbf{E}[K]$ grows linearly with T , while $\mathbf{E}[z^2]$ is bounded by a constant. This completes the proof. \square

A few points worth noting about Theorem 3.2: our lower bound on $\mathbf{E}[K]$ is actually $(\frac{\sigma}{4} - 2)T$ and σ must exceed 8 for this term to be positive and grow linearly with T . This is just

an easy way to handle the arbitrary integral rounding of Q_t . Intuitively, the algorithm typically cannot place orders with prices separated by less than a penny. Thus the volatility needs to be sufficient for integral changes to occur in the price series. If the unit of money could be made smaller, the loss due to rounding should diminish. Then σ in terms of the new unit increases linearly, while γ remains constant (Q_t becomes cQ_t for some scaling factor c). Thus for any constant σ , a sufficiently small granularity of prices allows us to apply Theorem 3.2. In fact, it is not difficult to see from the above analysis that the profit will grow linearly with time as long as the limiting variance of the process $\sigma^2/2\gamma$ is larger than (or comparable to) the granularity of bidding (normalized to 1). If this does not hold, then the algorithm will rarely get its orders executed, and will neither profit nor lose any significant amount.

Moreover, in the proof of Theorem 3.2, one may note that z is a normal variable with variance bounded by a constant, while the lower bound on K grows linearly with T , and occurs with high probability. Thus the profit expression in Theorem 3.2 not only holds in expectation (as is the case in Theorem 3.1 for general mean-reverting walks), but with high probability. Furthermore, for even the smallest of γ , Theorem 3.2 says that the profit is positive if T is large enough. Thus, *even when mean reversion is weak, a sufficiently long time horizon can make market making profitable.*

Finally, the profit expression of the OU process has a term $(\mu - Q_0)^2$. While we can treat it as a constant independent of the time horizon, we can also apply another trick to reduce this constant loss if Q_0 is far from μ . This is because Equation 2 tells us that the process converges exponentially fast towards μ — in time $t = \gamma^{-1} \log |Q_0 - \mu|$, $|\mathbf{E}[Q_t - \mu]|$ is down to 1, and $\mathbf{Var}[Q_0 - \mu]$ is less than σ^2/γ . Thus if the horizon T is large enough, then the market maker would like to simply sit out until time t (if allowed by market regulations), and then start applying our market making strategy.

3.2 The Schwartz Model

We now analyze another stochastic mean reversion model that has been studied in the finance literature and was studied by Schwartz [14]. The OU process assumes that the volatility of the price curve is a constant. Schwartz proposed a model where the volatility is a linear function of the price:

$$dQ_t = -\gamma Q_t (\ln Q_t - \ln \mu) dt + \sigma Q_t dW_t,$$

where μ is the long term mean price of the process, $\gamma < 1$, and $\sigma < 1$. Also assume that $Q_0 > 0$.

We shall show that the profitability of our market making algorithm for the Schwartz model is analogous to Theorem 3.2: $\mathbf{E}[K]$ grows linearly in T , while the expected loss due to liquidation $\mathbf{E}[z^2]$ is bounded by a constant, and hence the expected profit grows linearly in T . Applying Ito's lemma, Schwartz showed that $\log Q_t$ is an OU process, and so Q_t has a lognormal distribution, such that

$$\alpha_t = \mathbf{E}[\log Q_t] = (\ln \mu - \frac{\sigma^2}{2\gamma})(1 - e^{-\gamma t}) + Q_0 e^{-\gamma t},$$

$$\beta_t^2 = \mathbf{Var}[\log Q_t] = \frac{\sigma^2}{2\gamma}(1 - e^{-2\gamma t})$$

Then, by properties of lognormal distributions, we have

$$\mathbf{E}[Q_t] = e^{\alpha_t + \beta_t^2/2}, \quad \mathbf{Var}[Q_t] = (e^{\beta_t^2} - 1)e^{2\alpha_t + \beta_t^2}$$

Suppose the unit of price is small enough so that $\ln \mu > \frac{\sigma^2}{2\gamma}$ (again, this is essentially equivalent to choosing a finer granularity of placing orders). Note that shrinking the size of a unit step by a factor c leaves both σ and γ unchanged, but inflates μ by c . Since α_t and β_t are upper bounded by constants, so are $\mathbf{E}[Q_t]$ and $\mathbf{Var}[Q_t]$, and hence $\mathbf{E}[z^2]$ is bounded by a constant.

It remains to show that $\mathbf{E}[K] = \Omega(T)$. It suffices to show that $\mathbf{E}[|Q_{t+1} - Q_t|]$ is at least some constant. Note that Q_t is always positive, since it has a lognormal distribution. We shall show that $\mathbf{E}[|Q_{t+1} - Q_t| | Q_t]$ is at least some constant, for any positive Q_t . Since Q_t is a Markov process, this is equal to $\mathbf{E}[|Q'_1 - Q'_0| | Q'_0 = Q_t]$ for an identical but independent process Q'_t . Observe that $\alpha_1 \geq (\ln \mu - \frac{\sigma^2}{2\gamma})(1 - e^{-\gamma}) > 0$, if $Q_0 > 0$. Also, we have $\beta_1^2 > \sigma^2/4$. Thus

$$\mathbf{Var}[Q'_1] > e^{\sigma^2/4}(e^{\sigma^2/4} - 1) > \sigma^2 e^{\sigma^2/4}/4.$$

This shows that $\mathbf{Var}[Q'_1]$, given Q_0 , is lower bounded by a constant (that depends on σ). Since Q'_1 has a lognormal distribution, it follows that $\mathbf{E}[|Q'_1 - Q'_0|]$ is also lower bounded by a constant. This completes the proof for $\mathbf{E}[K] = \Omega(T)$.

4. TRADING FREQUENCY

Our price series model makes the assumption that the market maker can place fresh orders after every change in price. In practice, however, there are many traders, and each trade causes some change in price, and an individual trader cannot react immediately to every change. We thus consider a more general model where the market maker places fresh orders after every L steps. Let us consider the same market making algorithm as before, in this infrequent order model. Thus, for every i , at time iL the algorithm places orders around P_{iL} in a ladder fashion as before. These orders remain unchanged (or get executed if the requisite price is reached) until time $(i+1)L$, and then the algorithm cancels the unexecuted orders and places fresh orders. We say that L is the *trading frequency* of the algorithm.

The profit of the algorithm can no longer be captured succinctly as before. In particular, the profit is not exclusively determined by (nor can it be lower bounded by a function of) the prices $P_0, P_L, P_{2L} \dots P_{iL} \dots$ at which the algorithm refreshes its orders — it depends on the path taken within every interval of L steps and not just the net change within this interval. Still, some of our profit guarantees continue to hold qualitatively in this model. In particular, we simulate the OU process and run our algorithm on this process, to analyze how trading frequency affects the profit of the algorithm. We simulate an OU process with $\gamma = 0.1$, $\sigma = 1$ and the initial price equal to the long term mean.

First, we find that the profit still shows a trend of growing linearly with time, for different trading frequencies L that are still significantly smaller than the time horizon T . We simulate the algorithm with different time horizons T and different trading frequencies, and all of them show a strong linear growth (see Figure 2).

Also, the profit is expected to fall as the trading frequency increases (keeping time horizon fixed), since the number of trades executed will clearly decrease. We find that for a




Figure 2: Profit increases linearly with the time horizon, for different trading frequencies 1, 2, 5, 10, 20.




Figure 3: Mean profit decreases slowly with trading frequency (Horizon $T = 1000$).

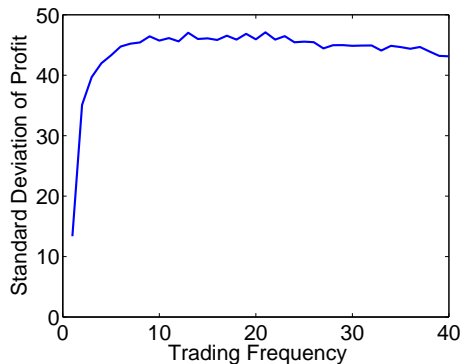


Figure 4: Standard deviation of profit increases quickly with trading frequency, then stabilizes (Horizon $T = 1000$).

large enough horizon ($T = 1000$), this is indeed the case, but the decrease in profit is quite slow, and even with trading frequency as high as 40, the expected profit is more than 80% of the expected profit with unit trading frequency (see Figure 3).

We computed average profit by simulating each setting 10000 times, to get a very narrow confidence interval. In fact, the standard deviation of the profit never exceeds 50 for any of our simulations, so the confidence interval (taken as 2σ divided by the square root of sample size) is less than 1, while the expected profit is much higher in all the cases.

The standard deviation in the profit itself goes up sharply as the trading frequency is increased from 1, but then quickly stabilizes (see Figure 4). The increase in variance of profit can perhaps be explained by the increase in variance of the number of shares that are liquidated at the end.

5. CONCLUSIONS

In this paper, we analyze the profitability of simple market making algorithms. Market making algorithms are a restricted class of trading algorithms, though there is no formal specification of the restrictions. Intuitively, the restriction is that a market maker has to always be present in the market, and offer prices that are close to the asset price. A future direction would be to put such a formal restriction, and try to design an optimal trading algorithm that satisfies the formal restrictions.

Acknowledgments

We give warm thanks to Yuriy Nevmyvaka for helpful comments and suggestions.

6. REFERENCES

- [1] C. Comerton-Forde, T. Hendershott, C. M. Jones, P. C. Moulton, and M. S. Seasholes. Time variation in liquidity: The role of market maker inventories and revenues. *Journal of Finance*, 65:295–331, 2010.
- [2] S. Das. A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance*, 5(2):169–180, 2005.
- [3] S. Das and M. Magdon-Ismael. Adapting to a market shock: Optimal sequential market-making. In *NIPS*, pages 361–368, 2008.
- [4] R. J. Elliott, J. van der Hoek, and W. P. Malcolm. Pairs trading. *Quantitative Finance*, 5(3):271–276, 2005.
- [5] M. B. Garman. Market microstructure. *Journal of Financial Economics*, 3(3):257–275, 1976.
- [6] E. Gatev, W. N. Goetzmann, and K. G. Rouwenhorst. Pairs trading: Performance of a relative-value arbitrage rule. *Review of Financial Studies*, 19(3):797–827, 2006.
- [7] G.E.Uhlenbeck and L.S.Ornstein. On the theory of Brownian motion. *Physical Review*, 36:823–841, 1930.
- [8] L. A. Gil-Alana. Mean reversion in the real exchange rates. *Economics Letters*, 69(3):285–288, 2000.
- [9] L. Glosten and P. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.
- [10] T. Hendershott and M. S. Seasholes. Market maker inventories and stock prices. *American Economic Review*, 97:210–214, 2007.
- [11] R. H. Litzenger and N. Rabinowitz. Backwardation in oil futures markets: Theory and empirical evidence. *The Journal of Finance*, 50(5):1517–1545, 1995.
- [12] S. Mudchanatongsuk, J. Primbs, and W. Wong. Optimal pairs trading: A stochastic control approach. *Proceedings of the American Control Conference*, pages 1035–1039, 2008.
- [13] R. Pindyck. The dynamics of commodity spot and futures markets: A primer. *Energy Journal*, 22(3):1–29, 2001.
- [14] E. S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997.