Market Microstructure: Time Series Analysis

Yuriy Nevmyvaka, Katia Sycara

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15232

Abstract

Market microstructure is a low-level high-frequency view of interaction among economic agents. In this article, we demonstrate the applicability of simple time-series models to market microstructure research. We show that a general autoregressive moving average framework is capable of capturing such microstructure effects as efficient price and mean-reverting bid-ask spread. This line of investigation can help us to both better understand the price formation process in financial markets and aid in construction of automated market makers and other trading agents. Simplicity of our models and normative approach differentiate us from other works in this area.

1. Introduction

The goal of this paper is to demonstrate how simple statistical techniques can be applied to market microstructure modeling. Market microstructure studies the price formation process in a given market. In financial literature, for example, it is often assumed that a security trades at a single price each day. In reality, however, price formation is a result of a complex negotiation process between buyers and sellers. Market microstructure examines the flow of different types of orders in the market place, evolution of the bid-ask spread, supply and demand for liquidity, intraday price patterns, etc. One practical application of market microstructure theory is in the field of market making – for-profit liquidity provision in a security. A market maker (or dealer) posts buy and sell quotes, thus giving other market participants a default counterparty, but also expecting to make a profit from this activity buy buying at lower prices than selling. For more details see [Nevmyvaka et al., 2003].

Most econometric models of market microstructure attempt to replicate the decision-making process of dealers and other market participants. This involves defining utility functions and then deriving conditions for an equilibrium. [Ho and Stoll, 1981], [Glosten and Milgrom, 1985], and [Seppi, 1997] are representative works of these "structured" models. In our paper, we will not analyze actions of specific economic actors, but instead just look at the data, searching for some underlying structure. Such investigation is useful in a number of ways: first, it gives us a better understanding of how financial market work and what the consequences of various events are. Second, this can be interpreted as an initial step towards predictive trading strategies, since we are essentially searching for short-term predictability in various microstructure variables, such as prices, quotes, spreads, volumes, etc. This knowledge can be used in constructing electronic market makers and other trading agents.

We are certainly not the first to suggest this approach – see [Roll, 1984] and [Madhavan and Smidt, 1993] among others. What sets us apart is our more practical approach: we use a simpler and faster model, and demonstrate its ability to capture several known microstructure effects. In our presentation, instead of a conceptual discussion, we rely heavily on numerical results. This paper is structured as follows: next section introduces our basic methodology – ARMA model, followed by a description of our dataset and its summary statistics. We then present resulting models for prices, spread, and introduce multivariate results. We conclude with a summary of our contributions.

2. Methodology

The purpose of this analysis is to investigate whether some useful information can be extracted from the historical data of stock transactions and limit order books without explicitly building a trading strategy around this information. There are essentially two approaches that can be used: first, we can take a time series of a single variable and try to find some underlying structure there (i.e. the series tends to revert to its historical mean with a certain lag); second, we can take several variables and look for dependencies among them (i.e. mid-spread leads transaction prices by a certain number of time steps). If any useful information is discovered through this type of statistical studies, it can then be profitably incorporated into a market making (or other trading) strategy. Below we outline the necessary steps – data collection, pre-processing, filtering, and model fitting, which can later be extended to more complex examples.

We tested the following hypotheses: (1) can past transaction prices help predict future transaction prices, (2) can spread size (together with past transaction prices) help predict future transaction prices, and (3) can the mid-spread help predict transaction prices. To run these experiments, we used frequently sampled MSFT (Microsoft Corp. stock) transaction data and order book evolution history during different time periods of one day. Essentially, we looked at three time series: transaction price, size of bid-ask spread, and midpoint of the bid-ask spread. We first fit the univariate ARMA (autoregressive moving average) model into each series searching for some underlying structure, and then used the spread size and the mid-spread to see if they can help model the transaction price.

All this analysis has been performed within SAS statistics software package – use [Delwiche and Slaughter, 2001] as a reference. For the detailed description of time series models see [Yafee and McGee, 2000]. Very briefly we introduce two basic stochastic processes. First, we make an assumption that adjacent entries in a time series are related to one another via some sort of a process, which can be described mathematically. There are many ways this can be done, but we are mostly interested in two types of time series models: moving average MA(q) and autoregressive process AR(p).

Under the one-step moving average process MA(1), the current output Y_t is influenced by a random innovation e_t plus the innovation from the previous time step:

$$Y_t = e_t - \theta_1 e_{t-1}.$$

The lag between t and t-1 needs not be one step, but can be any lag q or multiple lags. Another process can be such that the current output is determined by previous value plus some innovation:

$$Y_t = \phi_1 Y_{t-1} + e_t.$$

We call this an autoregressive process, which again can have an arbitrary lag p. (θ and ϕ are parameters for MA and AR processes respectively). These two processes put together form our main tool – process ARMA(p, q), which is simply a sum of the autoregressive and moving average components. Once again, notice that no microstructure variables enter this notation – in relaxed models, we deal only with numbers, regardless of where they came from.

We evaluate goodness-of-fit using three standard criteria: loglikelihood, AIC, and SBC. The first one is essentially a logarithm of the mean square error. AIC (Akaike Information Criterion) and SBC (Schwartz Bayesian Criterion) penalize mean square error with the number of features in the model. AIC = exp(2k/T)MSE and SBC = T^kMSE, where k is the number of features, T is the number of observations, and MSE is the mean square error.

Our experimental findings mostly confirm accepted principles from Finance Theory:

(1) markets do appear efficient (at least in a very liquid stock such as MSFT) showing little or no structure beyond white noise;

(2) size of bid-ask spread exhibits a fairly prominent AR(1) behavior in most cases;

(3) spread size does not help in transaction price forecasting;

(4) mid-spread is, in fact, useful for transaction price modeling, but only over extremely short time periods (3-15 seconds if that).

Although none of the above is revolutionary, these experiments highlight the power of multivariate ARMA models in market microstructure analysis. The exact same approach can help investigate more complex relationships: does volume misbalance signal upcoming price movement, does higher volatility lead to larger spreads, etc.

3. Dataset Description

The first dataset we use includes MSFT transaction prices collected from Island Electronic Communications Network on April 28th 2003 from 9:30 am to 4:00 pm. The price was sampled every 3 seconds, resulting in the total of 7653 data points. We then created a time variable starting at 9:30 am for the first observation, and incrementing it by 3 seconds for every subsequent observation. This main dataset (MSFT.DAY) serves as a base for the smaller time-of-the-day dependent datasets (see Table 1 below).

In order to concentrate on the short-term behavior of the transaction price, we selected three one-hour time periods during the day: beginning (10:30 am – 11:30 am), middle (12:30 pm – 1:30 pm), and end (2:30 pm – 3:30 pm). Note that we avoided using the opening and closing hour because, presumably, price behavior during these periods will be significantly different from "normal" rest-of-the-day behavior. The first 3 datasets created (MSFT.MORNING, MSFT.NOON, and MSFT.EVENING) are just the subsets of the master dataset in the indicated time periods. Since the price is sampled every 3 seconds, each of them contains 1,200 observations. In case if such sampling is too frequent, we also created 3 more datasets for the same time periods, but with the price sampled every 15 seconds. These datasets are called MSFT.MSHORT, MSFT.NSHORT, and MSFT.ESHORT and contain 240 observations each.

The second collection of data that we examined was a list of top bids and asks from the order book sampled at the same time as the transaction price. We used this data to create two more time series: the size of the bid ask spread, calculated as (Ask-Bid), and the mid-point of the spread: (Ask+Bid)/2. The later is often used in market microstructure theory as a proxy for the "true price" of a security. Then we went through the same steps as for the transaction price and ended up with 12 smaller time series: for both the size and the mid-spread we had MORNING, NOON, and EVENING periods sampled at 3 and 15 seconds each.

4. Basic Statistics and Stationarity

Tables 1 and 2 below summarize basic statistics for transaction prices and bid-ask spread respectively. The mid-spread dataset's statistics are essentially the same as those for transaction prices (Table 1), and thus are not reproduced here.

Name	Time	N	Mean	Min	Max
DAY	9:30-16:00	7653	25.7281	25.328	25.940
MORNING	10:30-11:30	1200	25.6798	25.551	25.770
MSHORT	10:30-11:30	240	25.6797	25.560	25.770
NOON	12:30-13:30	1200	25.8167	25.740	25.853
NSHORT	12:30-13:30	240	25.8169	25.740	25.853
EVENING	14:30-15:30	1200	25.8488	25.761	25.940
ESHORT	14:30-15:30	240	25.8489	25.761	25.936

Table 1. Basic statistics: prices

Name	Rate (sec)	N	Mean	STDev	Min	Max
DAY	3	7653	0.013639	0.0072	0.01	0.054
MORNING	3	1200	0.013039	0.00605	0.01	0.033
MSHORT	15	240	0.012846	0.0059	0.01	0.033
NOON	3	1200	0.01401	0.007384	0.01	0.054
EVENING	3	1200	0.013276	0.006776	0.01	0.043

Table 2. Basic statistics: spread

We can clearly observe from Table 2 the "U-shaped pattern" of the bid-ask spread, which is mentioned on many occasions: average spread is the largest in the middle of the day and is tighter in the morning and afternoon. The same holds for the maximum spread as well.

In order for the ARMA model to be applicable, the time series have to be stationary – in simpler terms, we had to remove the trend and render the volatility homoskedastic (roughly constant). All the series that involve prices (transaction or mid-spread) have unit root in them and must be first differenced. Dickey-Fuller tests in SAS ARIMA procedure prove that this is sufficient. It is much less clear, however, if taking logs of prices is in order to stabilize the series volatility. Results for several significance tests – log likelihood, AIC, and SBC – are presented in Table 3 for both regular prices and their logs.

Series	Log Likelihood	AIC	SBC
DAY	30621.56	-61231.12	-61189.46
Log	30608.10	-61204.19	-61162.54
MORNING	4660.74	-9309.48	-9278.94
Log	4660.77	-9309.54	-9279.01
MSHORT	752.908	-1493.82	-1472.96
Log	752.843	-1493.69	-1472.83
NOON	5288.68	-10565.36	-10534.82
Log	5288.57	-10565.14	-10534.60
NSHORT	870.565	-1729.13	-1708.27
Log	870.532	-1729.06	-1708.20
EVENING	4841.19	-9670.39	-9639.85
Log	4841.06	-9670.11	-9639.58
ESHORT	794.220	-1576.44	-1555.58
Log	794.137	-1576.27	-1555.42

Table 3. Stationarity tests

It appears that taking logs is not necessary for prices (transactions and mid-spread), but the difference is marginal.

Another issue is whether it is appropriate to work with actual prices, or should returns be used instead. The later approach is customary in financial literature, but may not matter for the kind of data we are using. To test which method is more appropriate, we initially fit all the ARMA models (see the following section) to transactions data using actual prices, and then replaced prices with log-returns, but left all the models parameters unchanged. Both approaches yielded the same results, so we chose to work with actual prices for other experiments.

We also determined that we needed to take logs (but not first difference) of the spread size for all time series. Spread size models had a significant intersection term, while prices did not. This can be attributed to first differencing of prices.

5. ARMA Models: Prices

We found it very challenging to fit an ARMA model to a time series of transaction prices, since they look very much as white noise. Surprisingly, however, when we extended the number of time periods to be examined by our model from 20 to 250 for 3 second series and to 50 for 15 seconds series, we found significant autoregressive terms that lag from 6 to 13 minutes:

Model	P	Q
MORNING	1,3,4	1,3,4
MSHORT	23,37	0
NOON	23,257,258	0
NSHORT	1,42,43	0
EVENING	6,117,118,230	0
ESHORT	1,46	0

Table 4. ARMA parameters: prices

While we are very much inclined to discard these results as nonsensical from the market microstructure point of view (a price at the next period depends on a price 10 minutes ago, but on nothing in between – sounds unlikely.), but these results have strong statistical support. Every single one of the parameters above is statistically significant (t-value is greater than 2), and both SBC and AIC are lower for the above models than for the base p=0, q=0 model; and, finally, the residuals are generally improving in most cases compared to the white noise model. Overall, if there is any underlying structure for transaction prices, it is almost certainly an autoregressive (as opposed to a moving average) relationship.

6. ARMA Models: Spread Size

Unlike transaction prices, spread size showed much more structure in correlograms: most of them look very similar to AR(1) model. AR(1) turns out first- or second-best model in AIC/SBC scoring, but some low-level (1 or 2) MA process seems to be present as well. Here are the parameters that we estimated:

Model	P	Q
MORNING	1	0
MSHORT	0	1
NOON	1	2
NSHORT	1	0
EVENING	1	1
ESHORT	0	1

Table 5. ARMA parameters: spread

This autoregressive behavior has a coherent explanation from the market microstructure point of view: as the spread narrows, it becomes cheaper for traders to “step over” the spread and transact immediately with outstanding limit orders; by definition, this will remove orders from the book and thus widen the spread. As the spread get wider, submitting market orders becomes more expensive, and traders resorts to posting limit orders inside the wide spread, which, in turn, shrinks the spread.

7. Multivariate ARMA Models

We next attempted to use the spread size and mid-spread as exogenous variables that help predict the transaction price. Whereas we did manage to find lags that make the spread size significant for the transaction price estimation, the new models’ SBC and AIC were always higher than the ones from the univariate model. Therefore, we reject our hypothesis that the spread size can be helpful for transaction price forecasting.

The mid-spread turns out to be a much more helpful variable especially when sampled every 3 seconds, which certainly is not surprising. We had to fit an ARMA model to the mid-spread series as well, again resulting in mostly AR models. After adding the mid-spread to the transaction price forecasting, we can conclude that in general, knowing the mid-spread at time t is useful for forecasting the transaction price at time t+1. SBCs and AICs are lower than without the exogenous variable, and lags are significant, but the residuals still leave a lot to desire in both cases. Does this finding have any practical significance? Not very likely, since one variable is leading the other one by an extremely short time period (plus lots of structure remains unexplained).

In all our experiments we obtained a vast amount of information describing significance of various coefficients, goodness of fit, behavior that still remains unexplained, some predictions, etc.; most of this data can also be plotted. But we are not reproducing all this numbers here because of the shear volume, and also because our primary goal is to demonstrate the process of finding out if there is some relationship between various microstructure variables and not to explicitly forecast stock prices or other variables.

8. Conclusion

We have described the basic idea behind the relaxed models and shown a simple application of time series techniques to market microstructure modeling. In our opinion, the main contribution of these experiments is the proof of applicability of multivariate ARMA models to the market microstructure research where we are dealing with discretely sampled data. We also found some structure in the spread size, which can mean that this variable is actually forecastable – a fact that can be used in creating automated dealers and other trading agents. And finally, while our efforts have confirmed that prices are hard (read impossible) to forecast, the same needs not be true for other microstructure variables.

Bibliography

1. Delwiche, L. and Slaughter, S., 2003, *The Little SAS Book: A Primer*, SAS Publishing.
2. Glosten, L., and Milgrom, P., 1985, *Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders*, Journal of Financial Economics 14.
3. Ho, T. and Stoll, H., 1981, *Optimal Dealer Pricing Under Transactions and Return Uncertainty*, Journal of Financial Economics 9.
4. Madhavan, A., Smidt, S., 1993, *An Analysis of Changes in Specialist Inventories and Quotations*, The Journal of Finance, Vol. 48, Dec 1993.
5. Nevmyvaka, Y., Sycara, K., and Seppi, D., 2003, *Electronic Market Making: Initial Investigation*, in Computational Intelligence in Economics and Finance, 2003.
6. Roll, R., 1984, *A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market*, Journal of Finance 39, 1127-1139.
7. Seppi, D., 1997, *Liquidity Provision with Limit Orders and Strategic Specialist*, The Review of Financial Studies, Vol. 10, Issue 1.
8. Yafee, R., and McGee, M., 2000, *Introduction to Time Series Analysis and Forecasting*, Academic Press.