Input and Output Standard Deviations in A Simulated Electronic Double Auction Market

2016 CEF 2016 Conference

Peter C. Anselmo

Department of Management

New Mexico Institute of Mining and Technology

15 February 2016

Abstract

Investor inputs into their decision-making processes as they allocate financial resources to different risky financial assets are an important driver of both theoretical and computational conclusions about the behavior of market prices. Many well-known and accepted financial theories assume efficient markets where information is dispersed across investors in a way that results in agreement about expected returns and the risks associated with those returns. This efficiency assumption results in analytically-tractable models and conclusions that to this day dominate conventional ideas about finance and financial economics.

On the other hand, relaxation of the efficiency assumption with respect to informational inputs to investor decision processes makes analytical results more difficult, but also creates an opportunity for the use of agent-based models – where investors/traders are the agents – to study price formation and price behavior in the context of differing levels of investor heterogeneity.

As a result, many agent-based models of electronic financial markets – or electronic stock markets – come with different assumptions (or, in the case of zero-intelligence traders, non-assumptions) about investor behavior that are used to generate simulated price time series that often are accurate representations of actual financial-market data. We will discuss our results in the context of the major alternative models of investor behavior that have been presented in the literature.

In this paper we develop the ideas outlined above with a focus on the return standard deviation input in a model of an electronic exchange market. Specifically, for this paper we present impacts on observed price time-series standard deviations in light of constant (and non-constant) values for other (other than the constant standard deviation) trader input variables. These impacts are studied in the context of the model of rational investor behavior in which simulated traders observe return distributions and make trading decisions based on the maximization of expected utility criterion.

1 Overview

The models used to represent trader agents in an efficient financial market have been shown to have an impact on whether model price outputs exhibit price and return behavior that correspond to well-known stylized facts present in real price dynamics (Lux and Alfaro, 2016). In this paper, we take a step back from much of the trader-modeling research of the last two decades by presenting the output of a model featuring only identical single-period expected utility maximizing traders with a uniform time horizon.

We also present details of our order-driven, double-auction market model in which prices are set by traders via their limit or market orders. In particular, the role of a liquidity-providing market maker is critical for generating transactions in this model. This role is explained below, and we note that without market-making of some type a model featuring only traders with identical information and decision-making processes will not exhibit price dynamics as all traders will submit identical orders much of the time. A liquidity provider therefore provides a mechanism for price changes, which in turn stimulates portfolio-rebalancing activity on the part of traders – and, therefore, some degree of dynamics.

This approach is motivated by a simple question: when we measure the standard deviation of a return time series[1] what is being measured with regard to the risk-measure inputs to the trader decision problem? If the trader/investor risk measure is expected return standard deviation, should the input standard deviation, which is unknown in reality but known to modelers, be the same as what is observed in the actual returns? The answer to this question may range from an emphatic “yes” to “it depends on the market structure” to “no”.

As we discuss below, our response is the middle one. The structure of the market itself, as well as the presence/non-presence of a liquidity provider, will have major impacts on market volatility given trader order flow. The liquidity-provider model we introduce below is trade, not price based, and represents an example of the importance of structural considerations alongside models of trader behavior as drivers of market dynamics.

We present a model where 50 trader agents make single-period portfolio rebalancing decisions involving a risky and riskless asset. The model features trader agents with beliefs that returns in each period (4 ticks in the model) are draws from the same independent and identically-distributed Gaussian distribution throughout the course of the simulation (Bachelier, 1900; Osborne, 1959). Agents optimize an exponential CARA utility function (Brock and Hommes, 1998; Brock, Hommes, and Waggener, 2005; Chiarella, Iori, and Perello, 2009; de Jong, Verschoor, and Zwinkels, 2010) over these non-changing data, and given the starting price for the simulation trade based on the last price and limit-book data available to all in order to balance their portfolios.

All trades in this model are based on a fundamental price identically determined by all identically risk-averse investors. Trading decisions are therefore portfolio-balancing decisions,

[1] Or pieces of a return time series via a sliding-window approach to sequential observations

as the existing market price and trader positions generate an exposure value that can be compared to the optimal, desired investment level at each iteration of the model.

The market is efficient as return and risk (standard deviation) information is equally available to all traders. The underlying assumption is that this information is an accurate depiction of the expected returns for the risky asset (Fama,1976). Information about constant return standard deviations that is efficiently distributed is also a foundation of long-lived financial ideas such as the Capital Asset Pricing Model (CAPM) (Sharpe, 1964; Lintner, 1965). An implication of the iid return assumption is unchanging standard deviation measurements that match investor inputs, and examination of that relationship (as noted above) motivates this work.

While the risk/return relationship has been tested in the context of CAPM and other models (e.g., Fama and French, 1993), these studies consider individual asset risk in the context of the overall market. This study represents an attempt to analyze the impacts of rational, efficiently-informed traders on the output standard deviation of the asset in question. We ignore overall market risk and portfolio effects and focus narrowly on a single risky and one riskless asset traded on an experimental electronic double-auction market. We compare the input standard deviation with respect to investor decision models, and the output standard deviation as measured using several moving-windows for empirically-generated asset returns.

In this model a key “agent” is the double auction market itself (described below and with flow charts in the Appendices). Along those lines, some relatively recent research into market dynamics includes the work of Rosu (2009), who modeled a limit-order book and developed analytical results based on volume of a single traded share. More recently, theoretical attention has been given to the specific role of the market-maker and market-maker inventory risks in the pricing of risk (Fournier and Jacobs, 2016). We note that there are numerous potential courses of action for any market maker at any point in the trading process, and the constant rules employed in this work are a naïve representation of that fact. However, we believe that the market-maker in this work is of interest, and that even a simplified order-driven market-maker as a the link between traders and prices is an important piece of the overall puzzle.

In the model presented here no learning process is modeled for trader agents, and we recognize that the fact that there is wide interest in agent learning models (e.g., Chen 2012) may make this paper of limited interest for many potential readers. Also, trader agents pursue one strategy: maximization of the expected utility (via maximization of the certainty equivalent) of their single risky asset/riskless asset portfolio. Trading occurs in the model as trader agents rebalance their portfolios based on changes in market prices – which in turn are driven by trader agent rebalancing decisions.

The trader agents in this model may therefore be seen as fundamental traders who never update information and never switch trading strategies. No mean-reversion coefficient is considered in the model. As may be seen in the simulation results presented below, this fundamental strategy does result in a self-fulfilling price time series of sorts, as the fundamental price implied by mean/variance information and agent risk aversion levels provides an anchor against large upward and (especially) downward price movements.

As mentioned above, the key to implementation of the model is the presence of a liquidity-provider agent, who is also able to trade as a "regular" expected-utility maximizing agent when there is no need to provide liquidity in the market. The market maker in this model provides liquidity by trading with an appropriate (defined below) limit or market order based on the condition of the book when a) a batch of 50 orders has been processed and b) no transaction has occurred after the order-processing is complete.

While different from a Walrasian market maker, this is similar in spirit to the market maker developed by Chiarella and He (2003) in that the market maker in our model acts at the end of a trading period. However, ours differs in that the market maker does not clear the market; rather the market maker trades with one order on either the limit order book or in the market-order queue. Many papers (Chiarella, He, Hommes, 2007; feature market makers that set prices separately from market-clearing activities, our market maker sets prices by transacting (in this case, in a pre-specified way described and diagrammed below). Other examples of price-setting market makers in which the market maker reacts to can be found in Shen and Starr (2002) He and Zheng (2010), Chiarella, He, and Zwinkels (2014).

In the last paper the market maker is modeled as reacting to aggregate demand at each time point in the model. In the model presented here, the market maker reacts in a similar way, but instead of clearing excess supply/demand the market maker will trade with the side of the market with more unfilled orders – and in our model there will be only one trade with the longest-waiting market order or the top order on the book. Thus only one order is filled, and traders then generate new orders as appropriate, and the next batch is processed.

Many simulations of limit-order books and double-auction markets generally (we define a double-auction market as one where there is both a limit-order book and market-order queues) have been focused on the impacts that learning agents have on price dynamics (LeBaron, 2002; Raberto and Cincotti, 2005; Anufriev and Panchenko, 2009). Our model differs from these listed according to the best-price/market order matching mechanism (LeBaron), the ability to queue market orders and the presence of a market maker (Raberto and Cincotti), and the combination of batch orders with a market-maker (a non-price setting market maker, as mentioned above) architecture (Anufriev and Panchenko).

The nature of the market maker (or liquidity provider) in this work illustrates the importance of the behavioral/market structure interaction discussed in Anufriev and Panchenko, and Botazzi, Dosi, and Rebesco (2005). While this work was not directly aimed at using modeling and simulation to study that interaction, the importance of the nature of the liquidity provider in this effort points to the need to study those interactions and their impacts on price dynamics – and therefore on trader behavior.

2 Model Protocol and Description

Following the outlines presented in Grimm, et. al. (2010) and Muller et al (2013), the simulation model is generally described in this section. We utilize the ODD + D framework of Muller et al because of the fact that this model features both trader agents and a liquidity-provider agent, and in fact the double-auction market itself can be considered an agent because many decisions about how orders are handled (prioritization of market orders and ways of handling very large volume orders are two examples of decisions made by the market liquididy-provider "agent") can be dynamically modeled using an agent-based approach.

Further, the model of the market maker has a major impact on the price time series generated via the model, and the description of that critical portion of the model – in this case, the model of the market maker represents the interface between the trader agents model and the logical protocol of the double-auction exchange mechanism – is included below.

2.1 Model Overview

The purpose of this study is to provide an agent-based-simulation approach to the question of whether the observed standard deviation of a return time series generated by a double-auction market is the same as the input standard deviation used by trader/agents as they buy/sell shares of a single asset based on a mean-variance expected-utility optimization model. Model development and analyses are motivated by the question of whether trader/investor standard deviation inputs generate similar (or identical) standard deviation outputs.

The model consists of two types of entities: trader/investors and the double-auction market. The market operates according to an unchanging set of rules, and there is a single-trader liquidity provider who has unique abilities to "observe" the state of the market and inject liquidity via a (probably) non-optimal transaction if no transaction has occurred during the simulated batch processing of 50 trader orders – including any orders generated using the expected-utility-maximization model by the liquidity provider.

Trader agents are characterized by a model of mean-variance rational investor certainty equivalent maximization (this is described below). Trader agents are all identically risk averse, and their investment horizons are another attribute that is the same across all traders. All trades also have the same access to the same inputs in the form of expected return (mean) and risk (standard deviation) information about the risky asset.

The overall exchange model can be seen in the flowchart in Appendix 1. Orders arrive in batches of 50 or less, with each order in the batch the result of the optimal portfolio-balancing decision made by each agent eligible to trade. Agents ineligible to trade are agents who have transacted less than 4 ticks prior to the current tick, and are waiting for the next time horizon to examine their positions and, perhaps, submit a new order. Agents who submit an order that is unfilled will leave that order in the queue or on the book until it is filled or until they submit a new order and cancel the existing order.

Trader agents are identical with regard to information and risk tolerance, so agent orders are processed sequentially without any sorting. In early stages of the simulation the limitation on the number of posted limit orders (15) may result in lower trader numbers having higher probabilities of early transactions, in the long run thetransaction counts for each agent are essentially identical.

2.2 Design Concepts

The central underlying concept is the model of trader agent behavior, which is the single-period expected utility maximization model that is the historical foundation for much theoretical financial economics. The central underlying assumptions are efficient markets and rational investors. Efficient markets means all traders have access to information about a same-sized window of historical prices, in this case in the form of historical expected return (ER) and standard deviation (σ) measures. Rational investors are described by the expected utility (and therefore certainty equivalent) maximization model shown below.

To test this idea in the context of rational trader agents, we implement the expected-utility-optimization based model of investor behavior in the context of a simulated electronic double-auction market. The simulation features a process in which fifty simulated traders submit orders which are then processed, prices are generated when a market ask is less than or equal to an arriving market bid (or the opposite) and/or a market order arrives, and then a new set of optimal order prices and volumes are developed by the fifty traders.

A simplified order-submission timeline is shown below. At time t, a trader generates a price/shares pair (L_{it}, S_{it}) that is the result of the decision process at that point in time. L_{it} is a limit price that may or may not result in a price observation P_t. The trader time horizon is m ticks (in this model, m=4 for all traders).

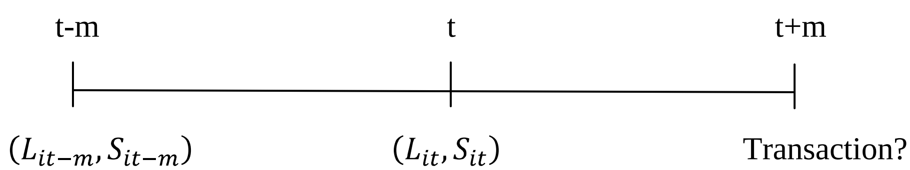

If a traders submits a market order, then only S_{it} is submitted to the batch list. The model has no memory, so previous trading success has no impact on the decision made at any point t in the process.

Order-generation is done in the context of a positive-return, risk-free alternative, whose rate remains constant throughout the simulated price-formation sequence. Investors maximize the certainty equivalent of their positions in risky and riskless assets in a normally-distributed returns context as follows:

$$\max_{\underline{V}} X|V_i = \left\{ \sum_{i=1}^{I} V_i(1 + ER_i) - \left(\frac{.5}{\theta} \left[\sum_{i=1}^{I} \left(2 \sum_{j \neq i} V_i V_j \rho_{ij} \sigma_i \sigma_j \right) \right] + V_f \left(1 + r_f\right) \right) \right\} \left(1 + r_f\right)^{-1}$$

ST

$$\textstyle\sum_i V_i + V_f = V_p$$

Note that this formulation encompasses the notion that the certainty equivalent X for a risky investment portfolio depends on the investment levels (Vs in the formulation) in each of the assets. An optimal vector of investment allocations V will result in a maximized portfolio certainty equivalent given expected returns ER, standard deviations σ, the correlation matrix with elements ρ_{ij}, the risk tolerance θ and risk-free rate r_f. The budget constraint V_p is shown as the constraint in the formulation.

Setting the derivative of the above formulation with respect to V_i equal to 0 and rearranging gives the I first-order conditions which lead to optimal investment levels V*:

$$V_i^* = \frac{\theta\left(ER_i - r_f\right)}{\sigma_i^2} - \sum_{j \neq i} V_j^* \rho_{ij} \frac{\sigma_j}{\sigma_i} \ \ \forall \ i$$

and

$$V_f^* = V_P - \sum_{i=1}^{n} V_i^*$$

In the above expression, V* denotes an optimal, certainty-equivalent-maximizing allocation level. It is easy to see the resulting formula for computing optimal risky-asset (and therefore riskless asset) allocation levels when there is only one risky and riskless asset.

We use the optimal risky asset allocation V* as the basis for the amount to buy or sell for each trader agent. These agents have to choose between three basic alternatives:

- Do nothing during the current period and remain with the status quo X_{it}
- Submit a market order with share volume S_{it}
- Submit a limit order with price and volume (L_{it}, S_{it})

The actual choice made will be based on the probability of filling a market order, the probability of filling a limit order, and the certainty equivalent values for each of the three alternatives.

We use the formulation above to compute certainty equivalents for market buy/sell orders and also for certainty equivalents for limit orders. Market expected returns are computed based on the last-observed market price vis-à-vis next-period expected return, and limit-order certainty equivalents are computed based on last-observed market price, next-period expected risky-asset return in the context of the limit price, and the probability of the limit order being filled as part of the next batch of processed orders.

Each trader has (in the case of the model presented here, identical) probability beliefs about the likelihood of a given limit order trading in the next period. These beliefs are based on either the prices on the order book, which is visible to all traders in our model, or the last (that is, most recent) price for the asset.

If we define q_t^L as the probability of execution for a limit order in the next (effectively current) period, then generally $\frac{dq^A}{dP^A} < 0$ for asks and $\frac{dq^B}{dP^B} > 0$ for bids. That is, probability of generating a trade is decreasing (increasing) as the ask (bid) price increases (decreases). In this formulation of the trader decision problem we assume that there are upper B_{it}^{Max} for bids and A_{it}^{Max} for asks, and lower A_{it}^{Min} for asks and B_{it}^{Min}price limits on execution probabilities. Thus, for a limit sell (buy), A_{it}^{Max} is the lowestpossible limit ask price at time t for investor i for which all asks greater than $A_{it}^{Max}$are judged to have an execution probability of 0 in the upcoming trading period. Similar logic applies to the lower ask limit and the low and high bid limits.

Execution-probability distribution bounds are computed based on a reference price Y_{it}^{R}, which is computed at each timestep (again, in this model identically, though that stipulation can be easily relaxed). Once Y_{it}^{R} is computed (see just below), limits on the execution price distribution are computed by all traders as follows:

$$A_{it}^{Max} = Y_{it}^{R}(1 + \sigma_{it})$$

$$A_{it}^{Min} = Y_{it}^{R}\delta_{it}^{-}$$

$$B_{it}^{Max} = Y_{it}^{R}\delta_{it}^{+}$$

$$B_{it}^{Min} = Y_{it}^{R}(1 + \sigma_{it})$$

where σ denotes the expected return standard deviation and δ is an adjustment factor that trader agents can use to subjectively adjust for the perceived short-term direction of the market price.

Computation of Y_{it}^{R} in the model based on the general state of the limit-order book. The book may be in one of 4 general states: limit orders on both sides, limit orders on either the buy or sell side but not on the other (a fairly common state of the book in the case of this model), and no orders on either side of the book. The simple calculations are done as follows (the prime notation on an ask or bid price denotes the best price on the book for the tick in question):

$$Y_{it}^{R} = .5(A_t' + B_t')$$

if there are both limit sell and buy orders on the book, and

$$Y_{it}^{R} = A_t'\alpha_{it}$$

if there are only asks on the book. Analogously,

$$Y_{it}^{R} = B_t'\beta_{it}$$

if there are only bids on the book, and

$$Y_{it}^{R} = P_{t-1}$$

if there are no posted orders on the limit book at time t.

While they may seem innocuous, α, β, and δ are critical parameters that are important for model tuning and stability. A discussion of the impacts of the selection of these variables (which, based on our experience, should all be valued close to 1) is beyond the scope of this paper.

The certainty equivalent of a limit order is generally defined as $X_t^L|Y_t^L$; we do not explicitly include share volume S in this simplified formulation. Recall that the certainty equivalent of the status-quo risky asset position is X_{0t}, which is the certainty equivalent of the portfolio is the limit order is not executed in the current period. The expected certainty equivalent for the asset position resulting form any limit order price is then

$$E(X_t^L) = (q_t^L|L_t^L)(X_t^L|L_t^L) + \left(1 - (q_t^L|L_t^L)\right)X_{0t}$$

Taking the derivative with respect to P (and omitting the t subscript) results in the first-order condition

$$\frac{dE(X^L)}{dL} = \frac{dq^L}{dL}(X^L|L)q^L\frac{d(X^L|L)}{dL} - \frac{dq^L}{dL}X_{0t}$$

For this work, we assume a uniform distribution bounded by the subjective price limits for the probability of execution given the limit price. Under these circumstances the the optimal ask $A_{it}^*$and bid B_{it}^* prices for all identical investors at time t based on solving $\frac{dE(X^L)}{dL} = 0$ go compute the optimal limit price on either side of the book at each iteration.

Note that the bid certainty equivalent is slightly more complicated than the ask formulation because of the dependence of expected returns on the bid price if the bid is executed.

For the market order choice, the expected certainty equivalent is defined as $X_t^M|V_t^*$, which is the certainty equivalent associated with the optimal asset valuation shown above in ( ). And as before, the status quo certainty equivalent at time t is X_t^0.

The expected limit order certainty equivalent $E(X^L)$ is compared with the expected market-order certainty equivalent $E(X^M)$ and X_0 so that the decision (P_t, S_t) at time t is defined as

$$(L_t, S_t) = Max\{X_t^0, E(X_t^M), E(X_t^L)\}$$

Where L_t is only specified as an optimal limit order price, and S_t is computed by each trader for market orders based on the current position and market price as:

$$S_t = \frac{P_0S_0 - V_t^*}{P_0}$$

which is the difference between the current position (P_0S_0) and the optimal position (V_t^*). In the case of a buy order, the terms in the numerator would be reversed.

Note that V_t^* will be the optimal currency amount invested in the risky asset given the (expected) certainty equivalent associated with the optimal strategy. For limit orders,

$$S_t = \frac{V_t^* - P_0S_0}{B_t^*}$$

for bids and

$$S_t = \frac{P_0 S_0 - V_t^*}{A_t^*}$$

for asks.

The homogeneous agents in this model interact only through the market, and there is no learning in this specialized case. As noted above, agents are homogeneous, though their transaction patterns may differ. This is because not all agents will have an order filled – or partially filled – at each tick. Thus as the simulation progresses some agents will be dormant and within their holding periods, while others will be seeking transactions.

Data collected in the simulation model include prices, volumes, buyer and seller identification numbers, bid and ask price depth and order depth, spreads, and market order depth. Also each trade type has a unique identification number (not relevant for this work).

2.2.1 Market Maker

The market maker in this model provides liquidity in the form of individual trades only in cases where a batch of orders has been processed with no resulting transaction This approach is driven by the modeling objective of facilitating transactions in a way that should have a minimal impact on the price stream and volatility.[2] In particular, the market-maker is designed to generate non-market-maker transactions without causing major directional changes in price.

Prior to assessing the market, the market maker is permitted to submit orders in each batch along with the other 49 trader agents. The model is structured so that the market maker cannot trade with herself, a stipulation that must be modeled with great care so that the price stream is not inadvertently forced into an upward or downward path that is not justified by trader agent order-submission behavior.

Instead of setting a price via a function based on limit-order book conditions, the market maker instead relies on a set of rules similar to the fuzzy rules developed for traders by Wang (2015). The IF/THEN rules followed by the market maker are as follows:

> IF Market Sell Orders are queued and no Market Buy Orders are queued THEN Transact with the top order and record the last market price and market-order volume.
>
> IF Market Buy Orders are queued and no Market Buy Orders are queued THEN transact with the top order and record the last market price and market-order volume.
>
> IF the number of queued Market Sell Orders is greater than the (nonzero) number of queued Market Buy Orders THEN transact with the top Market Sell Order ELSE transact with the top Market Buy Order.

[2]The assertion that this approach to market-making has minimal impact on simulation-generated prices has not been tested vis-à-vis other possible liquidity provision methods.

IF the (nonzero) number of Market Sell Orders and the (nonzero) number of Market Buy Orders are equal, generate a uniform random number and THEN transact with the top Market Sell Order if the random number is less than .5 ELSE transact with the top Market Buy Order.

IF No Market Sell Orders are queued and Limit Asks are on the Order Book AND the number of posted Limit Asks is greater than the number of posted Limit Buys THEN transact with the best ask at the limit price and volume.

IF No Market Buy Orders are queued and Limit Buy orders are posted on the book AND the number of posted Limit Buy orders is greater than the number of posted Limit Sell orders THEN transact with the top limit buy order at the limit price and volume.

IF no Market Orders are in either queue and the number of posted Limit Asks is greater than the number of posted Limit Buys THEN transact with the best Limit Ask at the limit price and volume.

IF no Market Orders are in either queue and the number of posted Limit Bids is greater than the number of posted Limit Asks THEN transact with the best Limit Buy at the limit price and volume.

If no Market Orders are in either queue and the number of posted Limit Bids and Limit Asks is equal, generate a random draw from a uniform random distribution and THEN transact with the Limit Bid if the random number is < .5, ELSE transact with the Limit Ask.

In this special case of certainty-equivalent-maximizing trader agents, existing limit and market orders are deleted prior to the generation of new limit orders for the batch. In the case of a trader generating a new market order for the batch, existing limit orders are deleted (and the book is adjusted). In the case of a trader with an existing queued market order generating a new market order for the batch, the new order is deleted in order to maintain the queue position of the existing order – provided that it is on the same side of the book. Otherwise the existing market order is deleted.

Because of this protocol, the sequence of logical statements driving the market maker has no impact on price time series generated by the model. This is because the identical traders often generate identical orders that are identically on the same side of the book. In the case where traders are heterogeneous, the market maker's liquidity-provision function becomes slightly more complicated given the goal of limited intervention in and/or causation of directional price movements.

As noted above, the market maker is also one of the 50 traders in the model. As such, the market maker will manage her share inventory via submission of conventional orders – orders that, despite the market maker's identical input mean-variance information and risk tolerance, will likely differ from orders submitted by the other 49 traders because of more extreme inventory positions caused by market making.

2.3 Details

The model is currently implemented in *VBA Excel*, and one simulation run consumes around 90 minutes on a multi-core machine. The code is not publicly-accessible at the moment[3] In addition to critical variables held as constants throughout the simulation runs (see Table 1 below), the critical initial state that matters is the price P_0 at the first trader decision point. As long as $P_0 \neq$ (the optimal market price – in this case, 49.90) given share holdings and risk preferences for all traders, discrepancies between optimal asset allocations from V* and the market-price driven allocations will result in trading. If the simulation is initialized with $P_0 = 49.90$, traders will not submit orders absent additional information – which is not a feature of this simplified model.

Input data, as noted above, were held constant throughout the simulation. Input data for traders are expected return, std dev, risk tolerance, risk free rate, which are constants. Other input data for traders are limit order best prices if they are available.

Table 1. Values of constants used in the simulations.

Variable	Notation	Value in Model
Expected Period Return	ER	.025
Expected R Standard Deviation	σ	.05
Risk Free Rate	r_f	.00005
Risk Tolerance Parameter	θ	50,000
Trader Investment Time Horizon (Ticks)	m	4
Simulation Initialization Price	P_0	Variable; Not the Fundamental Price
Number of Trader Agents	N	50
Total Number of Outstanding Shares	**S**	500,000
Initial Trader Share Allocations	S_{i0}	10,000
Trader Budgets	V_p	Unconstrained
Bid Distribution Reference Price Factor	β_{it}	1.05
Ask Distribution Reference Price Factor	α_{it}	.95
Bid Distribution Trader Upper Bound Factor	δ_{it}^{+}	1.005
Ask DistributionTrader Lower Bound Factor	δ_{it}^{-}	.995

3. Simulation Results

This modeling and simulation effort was motivated by a simple question: are observed return standard deviations comparable to (in reality, of course, unobservable) standard deviations which are inputs to investor optimization decisions? In reality, of course, asset expected return standard deviations as risk-measure inputs are not likely to be observable. These measures can be specified in an experimental simulation model and compared with returns on price time-series outputs.

[3] Please contact the author for information about accessing the program.

Thus the primary purpose of this modeling and simulation effort was simple: compare observed return standard deviations with inputs to the rational, single-period trader certainty-equivalent optimization models. As noted by Anufriev and Panchenko (2009), market structure and order-handling logic are important for understanding price dynamics, and the results presented in this section should be considered with that point in mind. To illustrate this point: for the data reported here, 98.5% of all transactions were with the market maker.

Major results from this simulation effort are as follows:

- Return time series do not exhibit patterns in accordance with stylized facts. Specifically, when the time horizon used to measure returns is matched to the time horizon modeled for trader agents, there are no return volatility clusters (Figure 3 below).
- For tick-to-tick returns (m=1) and 4-tick returns (m=4) the near-constant volatility generated by the price process is smaller than the initial input standard deviation of .05. As m increases volatility becomes non-constant and (in the cases where m=500 and m=1000) larger than the input value.

Periodic returns given tick-interval m were assessed as:

$$R_{t|m} = P_t / P_{t-m} - 1$$

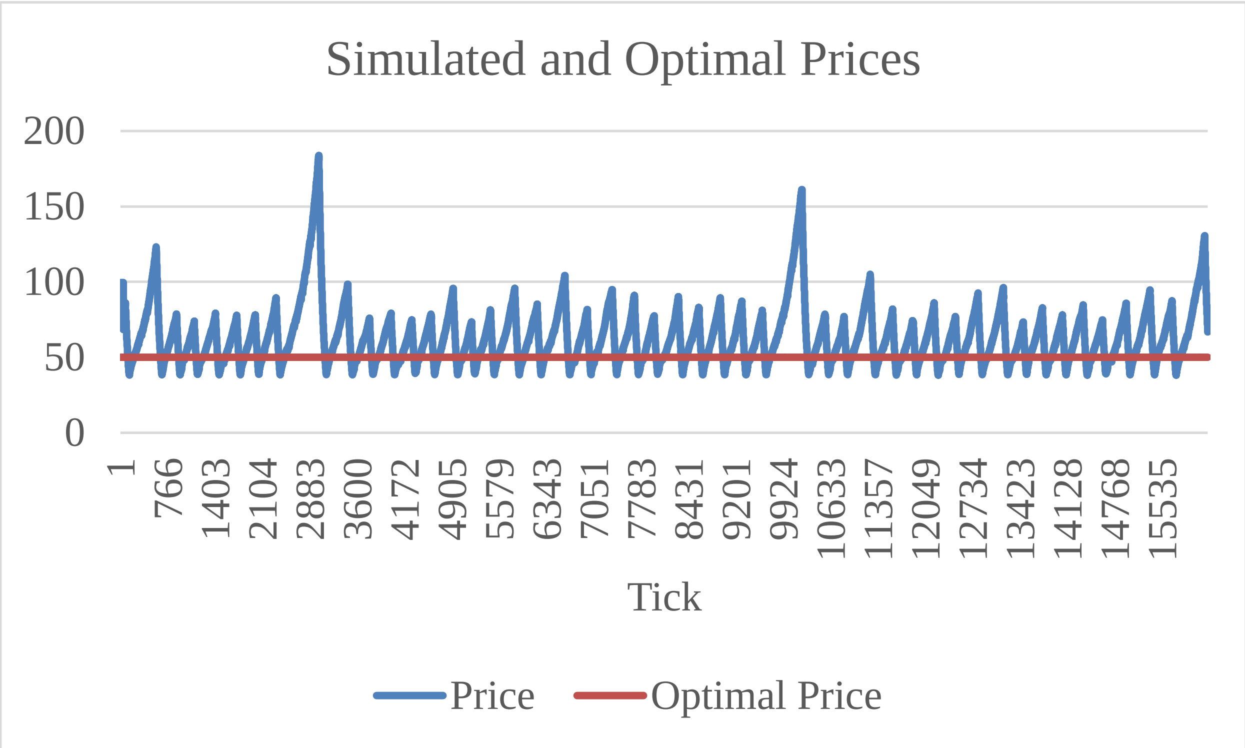

Figure 1. Simulated and Optimal Prices when n, the investment horizon, m=4. The optimal price of 49.90 does not change over time because of the homogeneous traders and investment-decision inputs.

An example of one of the generated price time series is shown above in Figure 1. The observed price and implied optimal price[4] are shown. One implication of the price stream is that this model may provide experimental evidence for the existence of a fundamental market price (though perhaps not a fundamental return). However, this idea must be considered in the context of the trader model and the way the experimental market is structured. If there is no liquidity provider in a market with identical traders, there will be no fundamental price that is based on aggregate investor preferences and the supply of shares.

The return time series for one of the simulation runs when m=4 is shown below in Figure 2. Note the apparent absence of volatility clusters, a stylized fact normally apparent in financial-return data (Lux, Alfaro 2016). Given the price series shown in Figure 1, the returns in Figure 2 are not surprising, and lead to the question of whether the notion of a fundamental price and the stylized fact of volatility clusters in observed returns for many auction-market-traded assets are compatible.

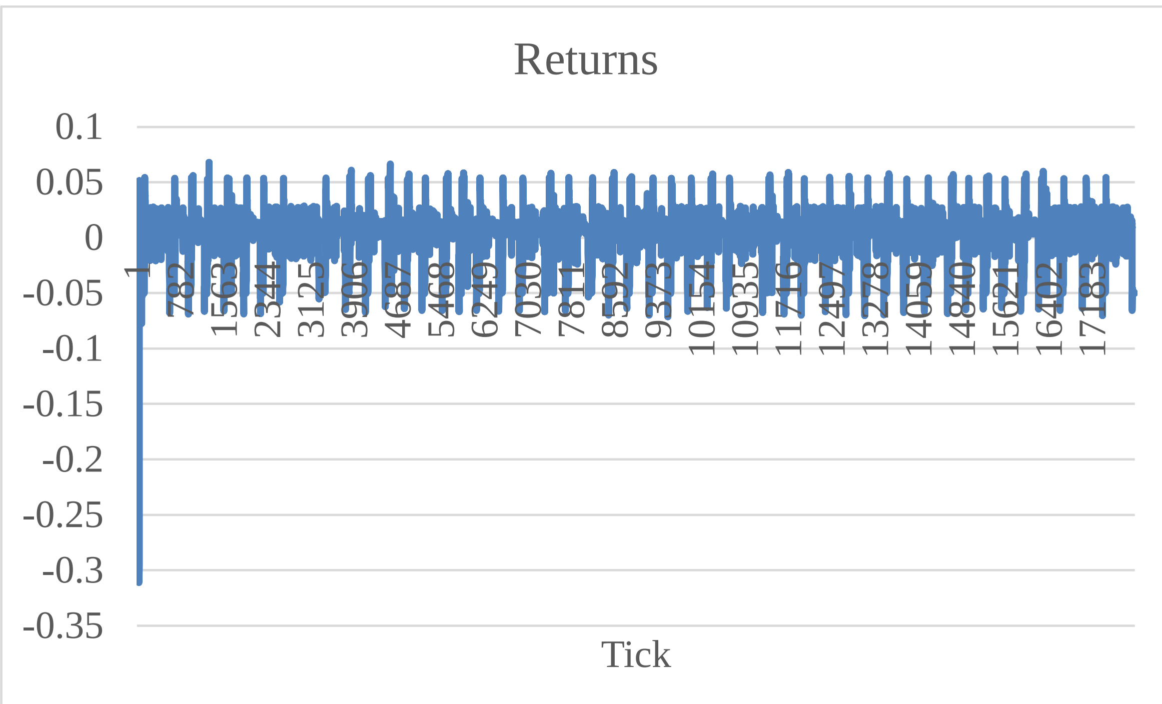

Figure 2. Simulated 4-tick price returns in the model. The initial low returns are associated with the initial randomly-generated price, as the simulation evolved into a steady state.

The average return over the course of the simulation was .00001 (m=1) and .0002 (m=4), which also represent departures from the input expected return of .025.

Tick-by-tick standard deviations for the price time series shown in Figure 1 are presented below for m = 1 and m = 4 are shown below in Figure 3. The overall standard deviations for all return

[4]The optimal price is generated based on the solution to the trader optimization formula in ( ) and, since there are 50 identical traders, equal share allotments. The same number ($49.90 given the input variables in Table 1) can be computed by summing the aggregate optimal valuations and dividing by the total number of shares.

observations were .0085 (m=1) and .024 (m=4). Recall that the input standard deviation was .05 for all investors. Standard deviations as shown in Figure 3 were measured using a rolling-window method. Window sizes were 1000 and 1500 observations.

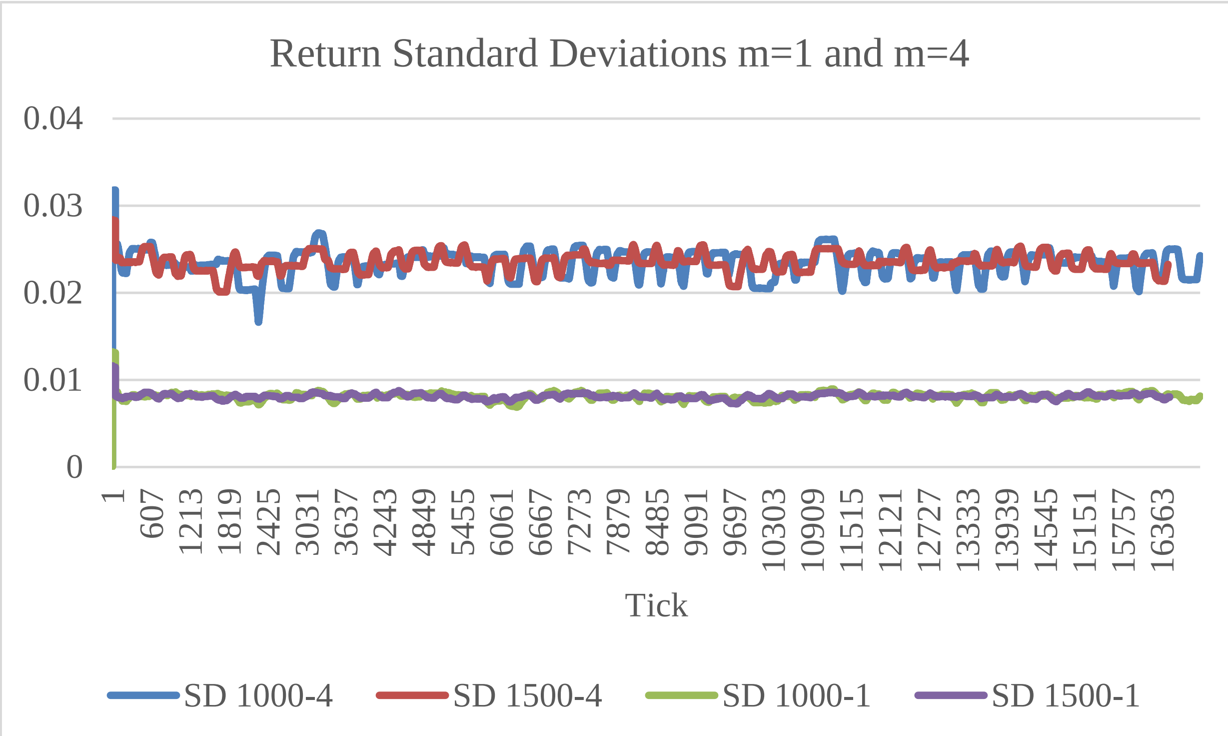

Figure3. Running standard deviation measurements for the double-auction market model with input standard deviation = .05 for all investors. Two sets of standard deviation measurements are shown for values m = 1 and m =4. Within each set are two ranges of standard deviations measured over running intervals of 1000 and 1500. The legend is interpreted as follows: SD 1000-1 is the graph of running standard deviations for 1000 return observations when m=1, etc.

In order to examine the impact of lengthening the return-measurement period on the rolling-window standard deviations, returns were computed for m=500 and m=1000 and standard deviations were computed. Graphical results are shown in Figure 4. Since these results are from the price time series generated by modeled traders with a 4-tick horizon, Figure 4 should be interpreted with some care.

However, return standard deviations that increase in m are consistent with observations of actual data. Return observations for real prices made in the context of an assumed rational model for investor behavior might be done in the case of multiple m values with an eye toward stability providing a notion of an underlying investor time horizon.

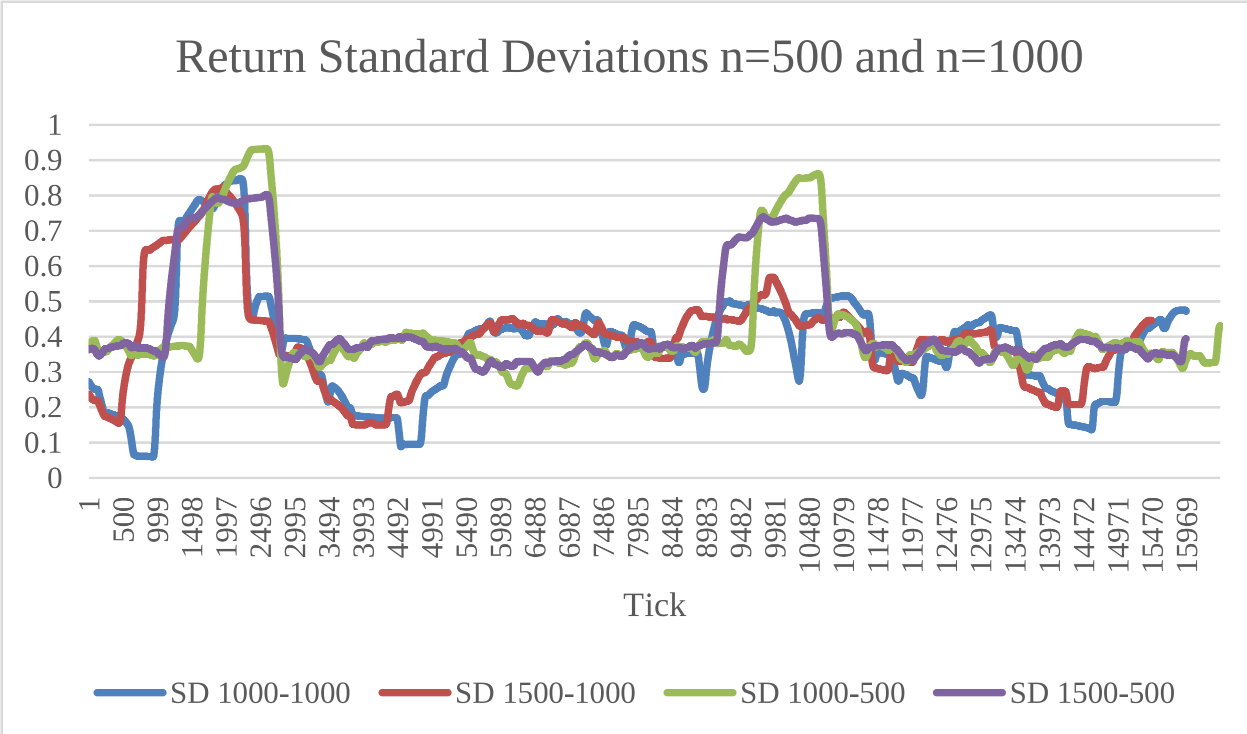

Figure 4. Running standard deviation measurements for the double-auction market model with input standard deviation = .05 for all investors. Two sets of standard deviation measurements are shown for values n= 500 and n=1000. Within each set are two ranges of standard deviations measured over running intervals of 1000 and 1500. The legend should be interpreted as follows: SD 1000-500 is the running standard deviation measure of size 1000 for the case where n=500 for return computations.

The results presented in this section are from one of 10 individual simulation runs, and while the numbers are unique, the overall patterns and conclusions (especially the observed standard deviation < .05 for m=4) are the same. The initial price P_0 was, as noted above, randomly generated. Regardless of the initial value for P_0 – with the exception of the case where P_0 = 49.90 – the price stream quickly converged to a stable process similar to Figure 1.

4. Conclusion

We have presented summary results for a double-auction model featuring homogeneous fundamental-price investors. The goal was to examine the standard-deviations of price returns given the (constant) standard-deviation inputs to investor decision models. As noted above, the nearly-constant output standard deviations for return observations over intervals matching hypothetical trader time horizions, while constant, do not match the input standard deviations.

This result, and any result from a model of a double-auction financial market with market queuing, must be considered in the context of modeling decisions about order-handling priorities and processing and, particularly, how the liquidity provider is modeled. The liquidity provider in this model, while very active (this agent was involved in approximately 98.5% of all transactions), is not a price-setter and is only a trade-facilitator. For future work a better understanding of how market-making actually happens, particularly in a dynamic context, would be very advantageous.

In addition to the model and design-driven answer to the motivating question for this work, there are some additional observations generated by this effort.

First, trading and volatility can happen in a market with identical fundamental-price traders, but only if there is a point where prices cause exposures to differ from optimal values to the point where investors will attempt to trade to re-balance their portfolios. Such orders will not result in trades in this situation, though, unless a liquidity-providing market maker is present.

Second, market structure (e.g., market queue inclusion) and market-maker agent decision models are extremely important. As of now, a very important market-maker design issue seems to be the choice about volume handled when providing liquidity. For us it's the volume requested by the appropriate opposite order, and while that makes sense to us there are many other ways to approach this point.

Third, the model by which simulated traders develop order-filling probabilities that is introduced above should be tested further and, if possible, verified with real market data. While the latter would be a difficult task, inferring market limits on upper and lower bounds of the price distribution for order-filling from real data would be a good step forward in the development of aucton-market models.

Finally, while there appears to be no correspondence between input and output standard deviations in the context of this model, it does appear to be possible to generate a price stream with an output return standard deviation that is a close match to the input standard deviation in the model by modifying the market structure, particularly the actions of the liquidity provider. However, actual input standard deviations are not observable, and situations where observed standard deviations are used as decision-model inputs may need to be closely examined.

APPENDIX 1. EXCHANGE ENGINE FLOWCHART

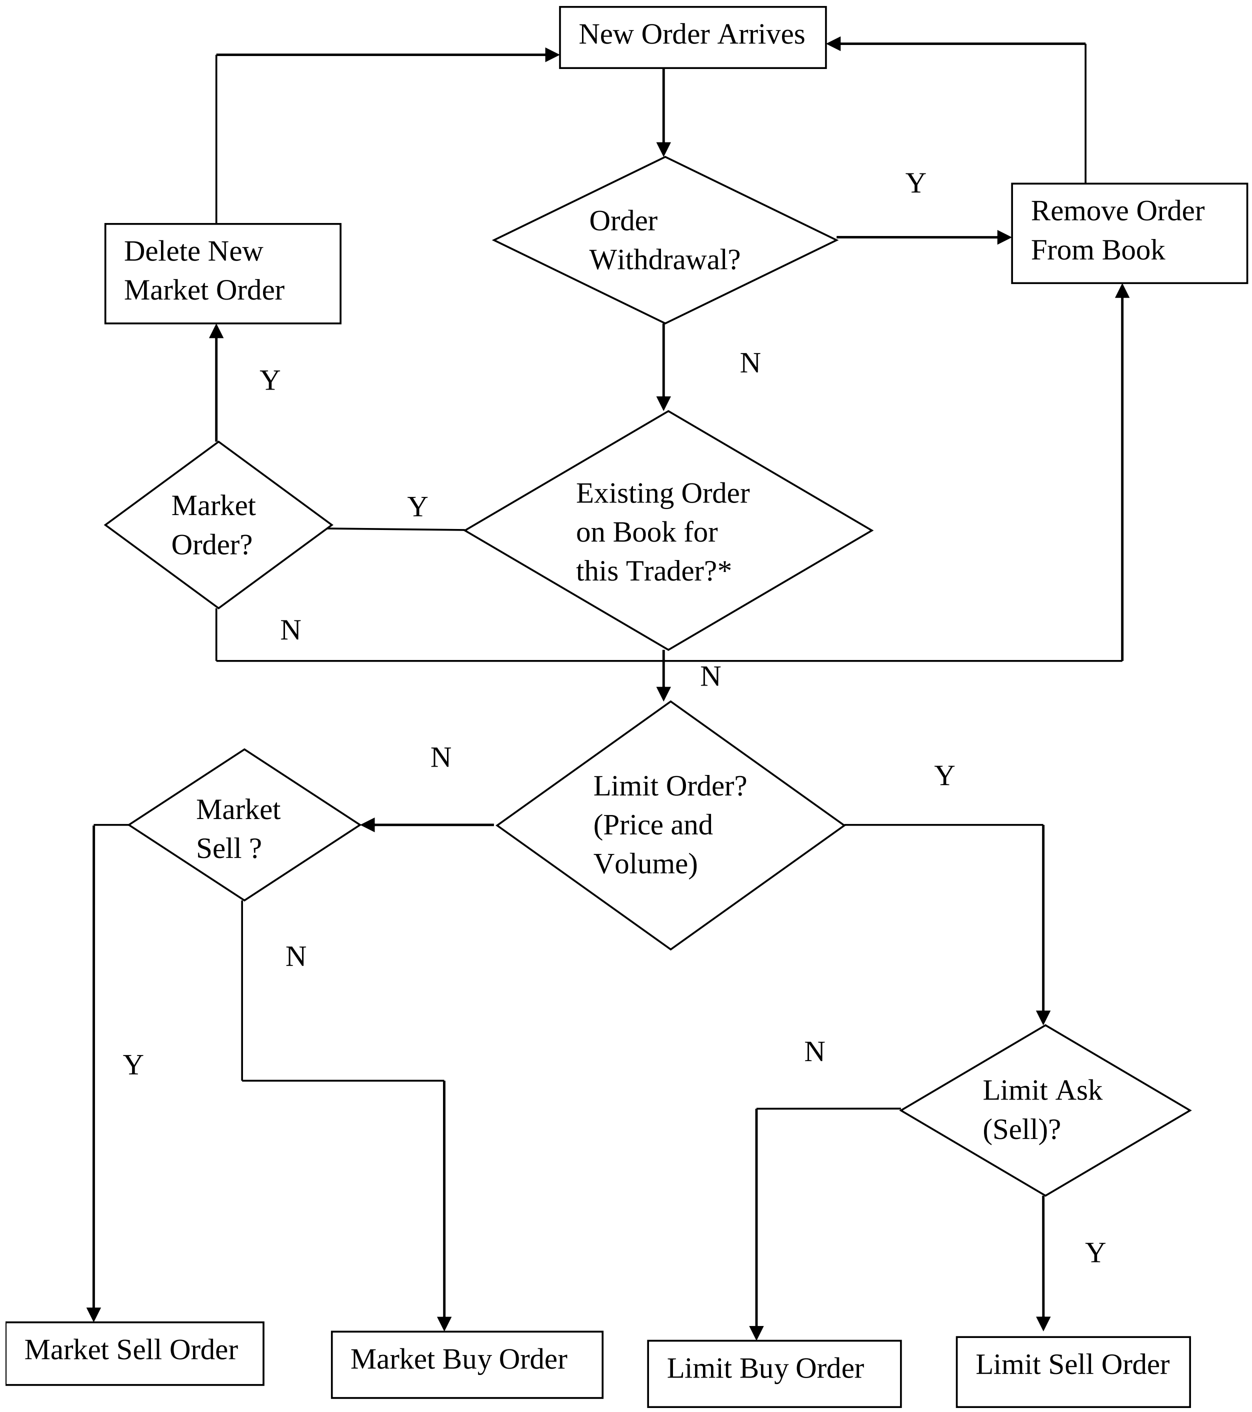

*This requirement can be relaxed or modified depending on modeler preference.

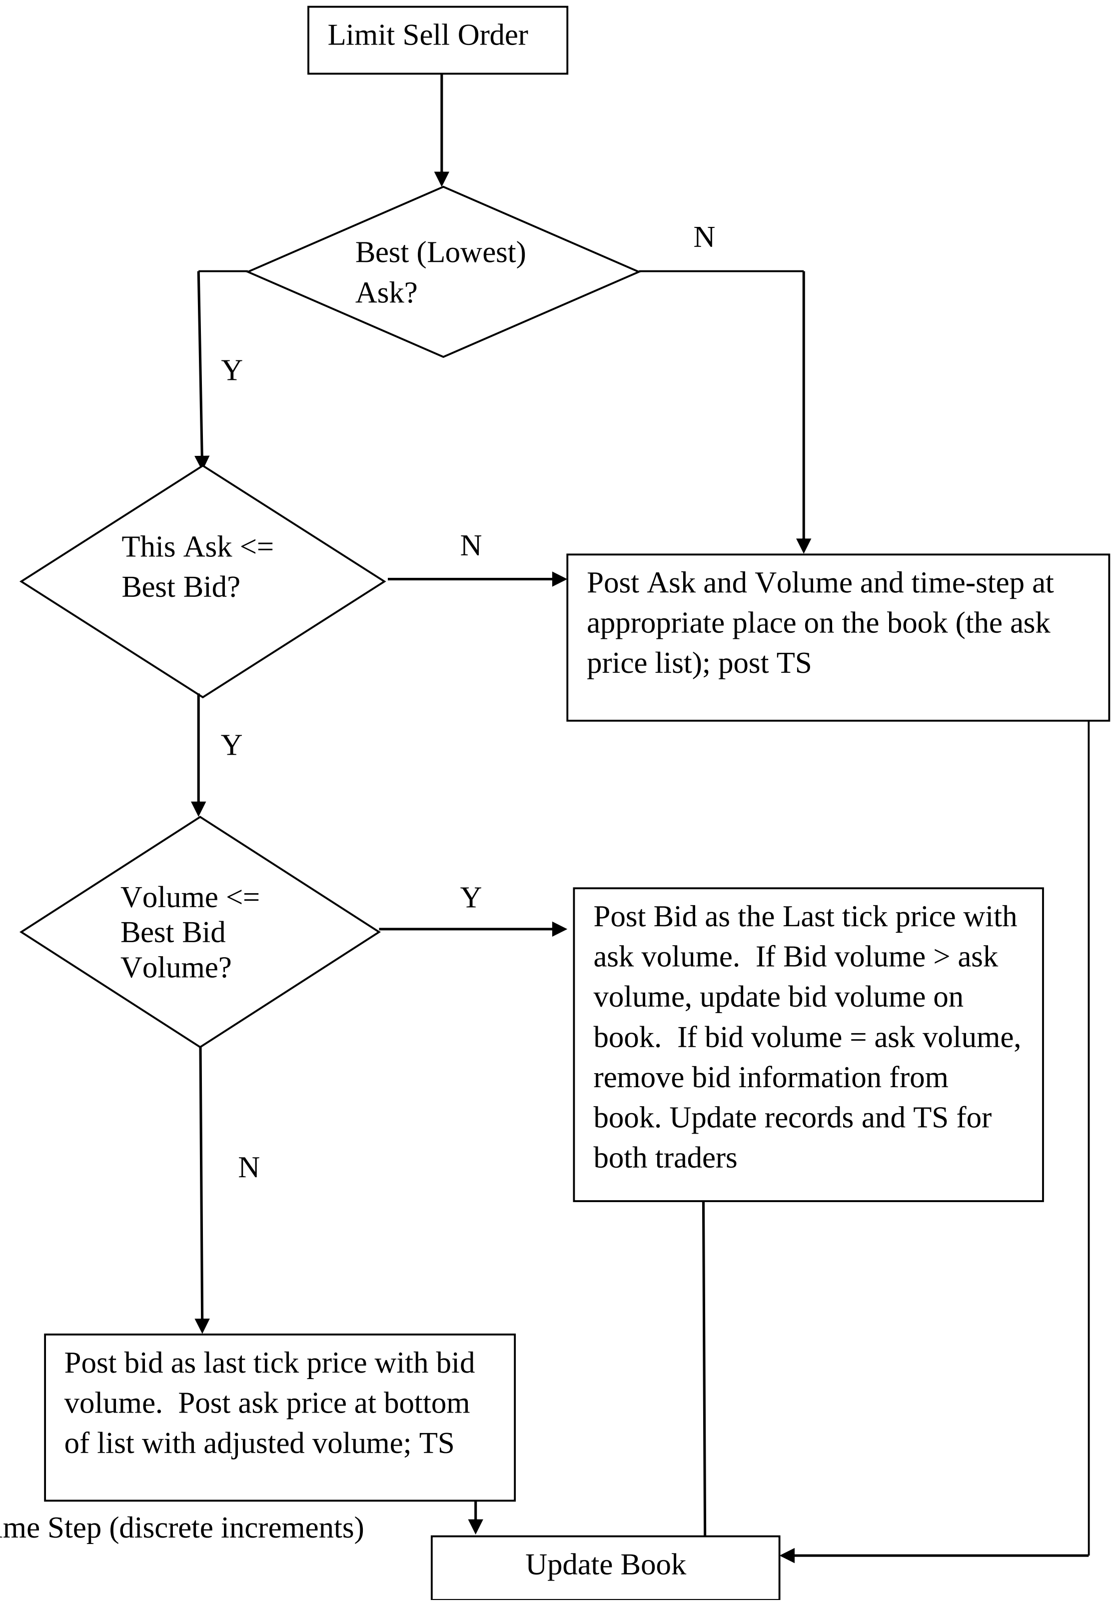

** TS = Time Step (discrete increments)

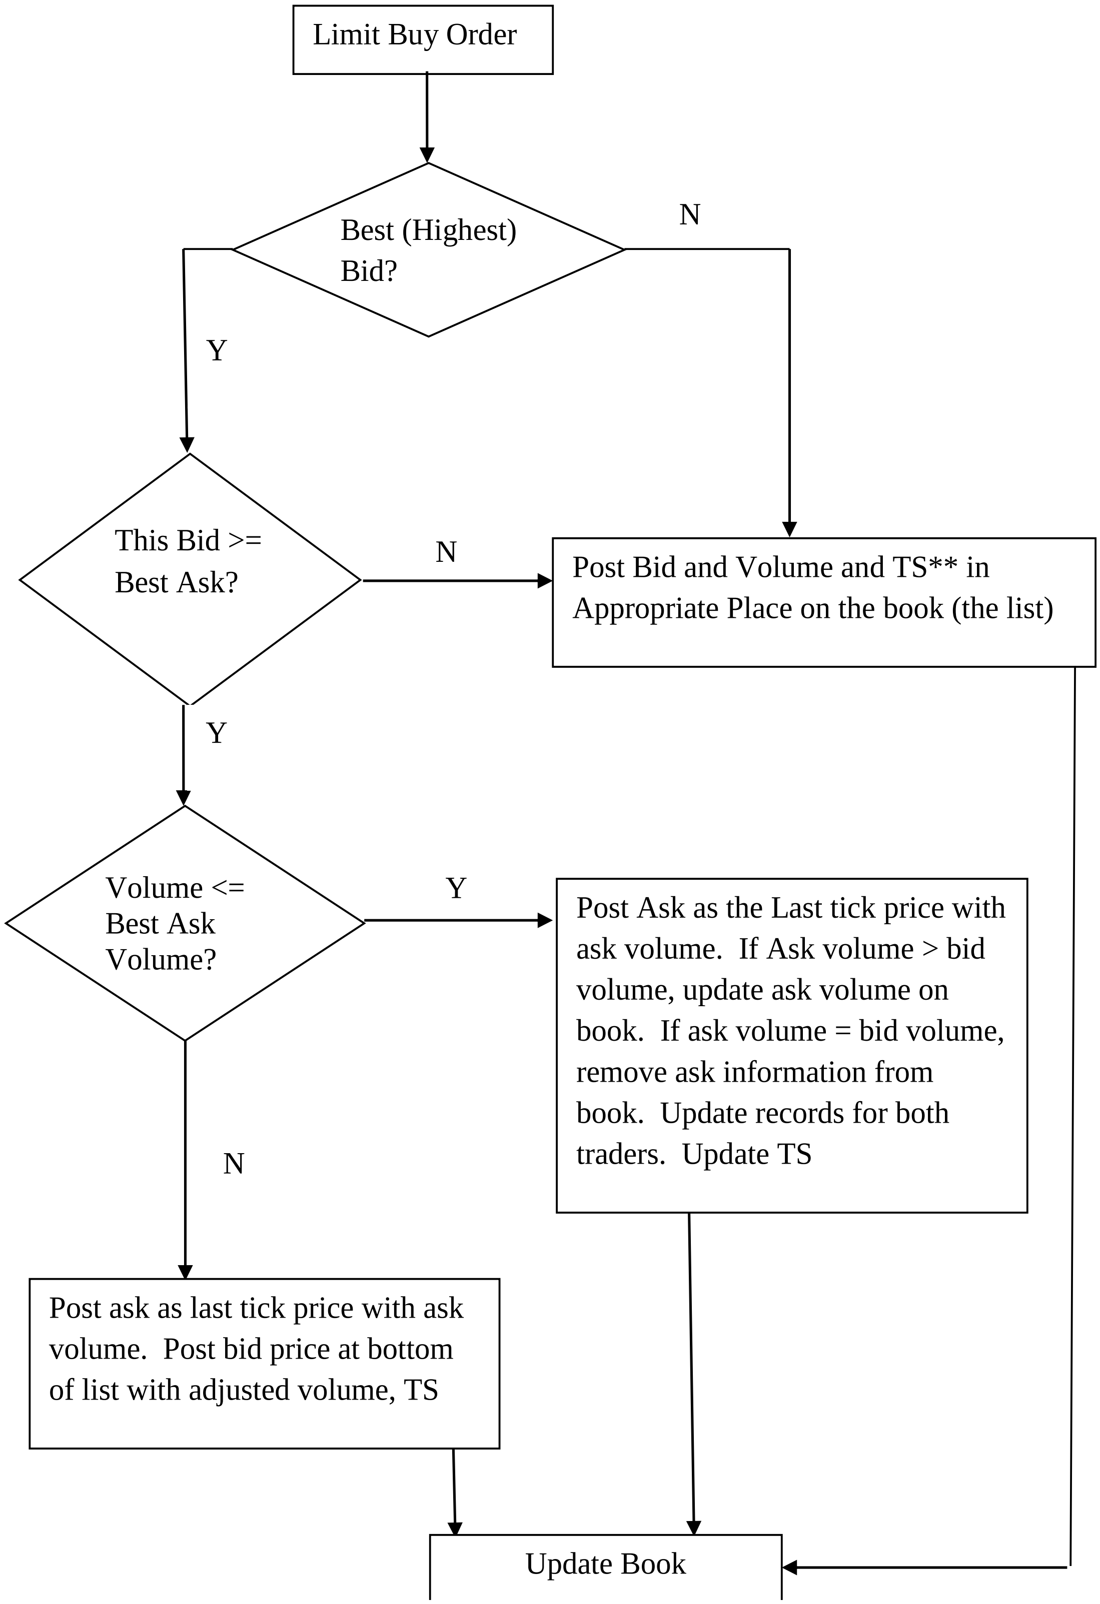
Limit Buy Order
Best (Highest) Bid?
N
Y
This Bid >= Best Ask?
N
Post Bid and Volume and TS** in Appropriate Place on the book (the list)
Y
Volume <= Best Ask Volume?
Y
Post Ask as the Last tick price with ask volume. If Ask volume > bid volume, update ask volume on book. If ask volume = bid volume, remove ask information from book. Update records for both traders. Update TS
N
Post ask as last tick price with ask volume. Post bid price at bottom of list with adjusted volume, TS
Update Book

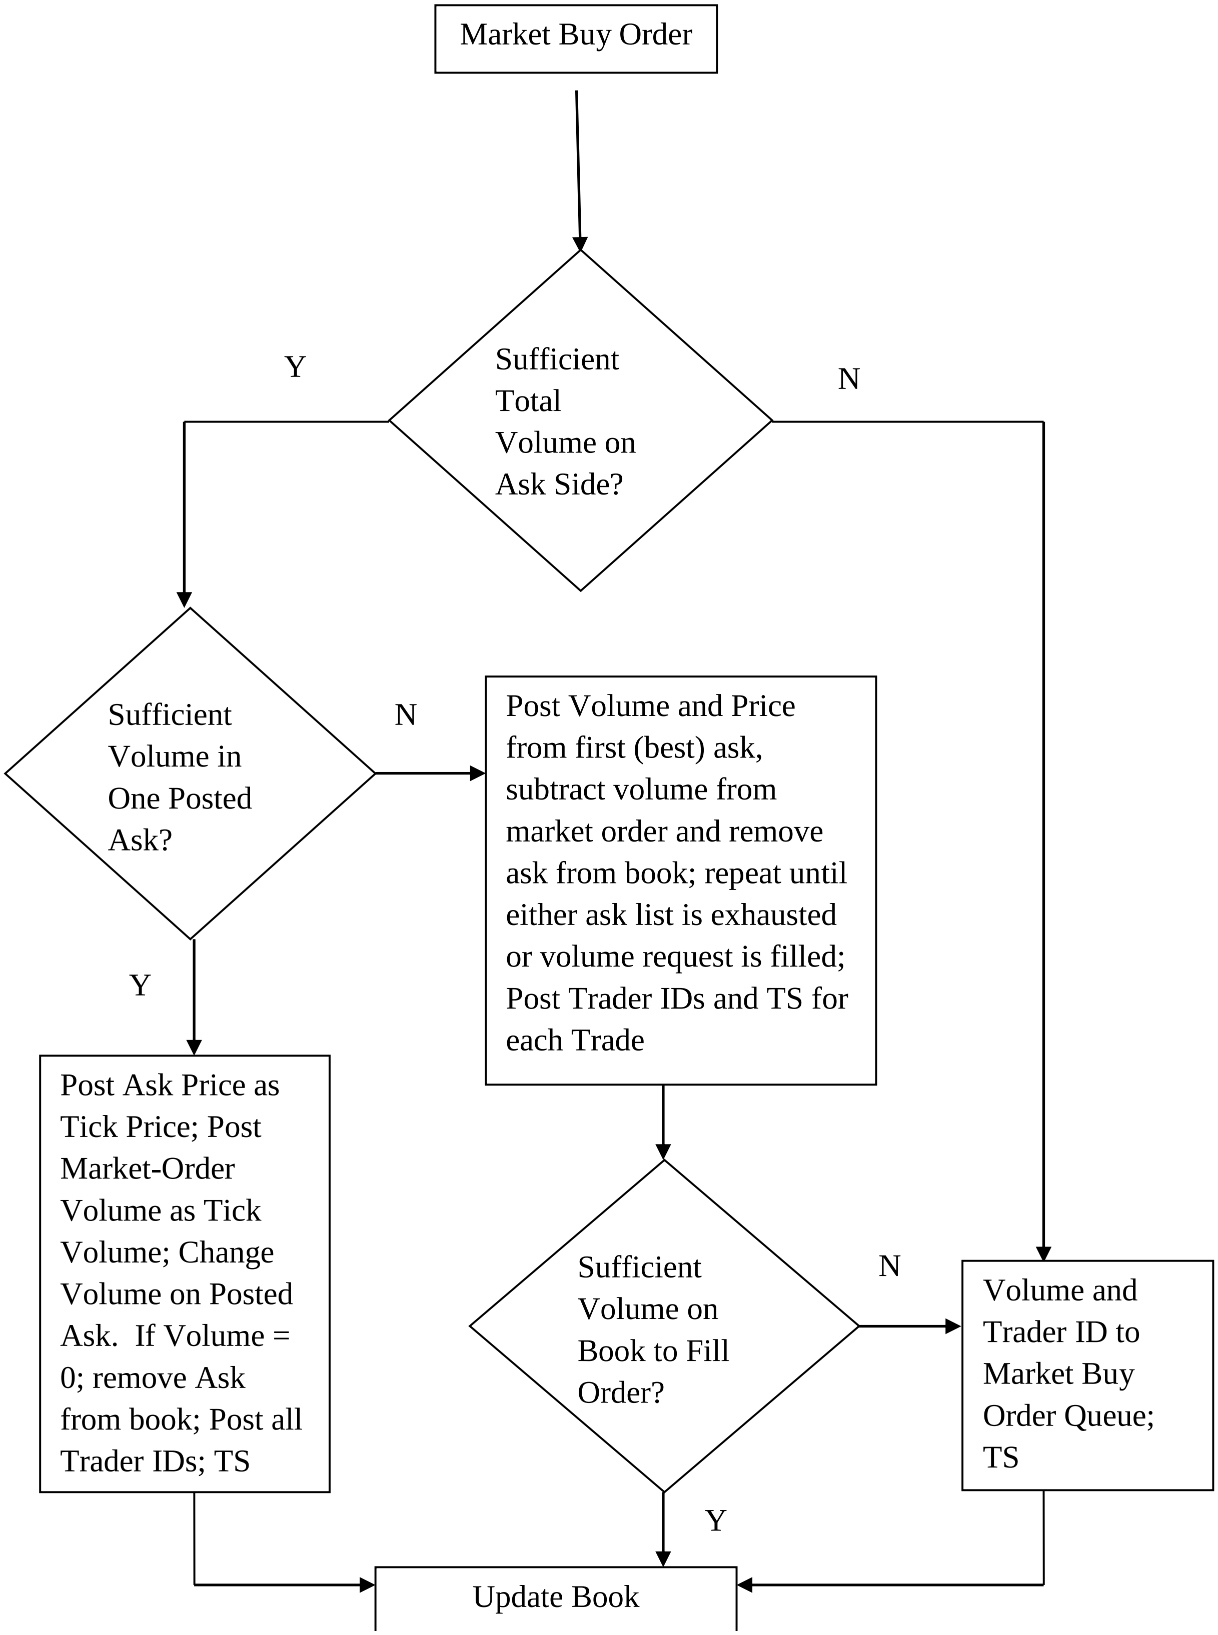
Market Buy Order
Sufficient Total Volume on Ask Side?
Y
N
Sufficient Volume in One Posted Ask?
N
Post Volume and Price from first (best) ask, subtract volume from market order and remove ask from book; repeat until either ask list is exhausted or volume request is filled; Post Trader IDs and TS for each Trade
Y
Post Ask Price as Tick Price; Post Market-Order Volume as Tick Volume; Change Volume on Posted Ask. If Volume = 0; remove Ask from book; Post all Trader IDs; TS
Sufficient Volume on Book to Fill Order?
N
Volume and Trader ID to Market Buy Order Queue; TS
Y
Update Book

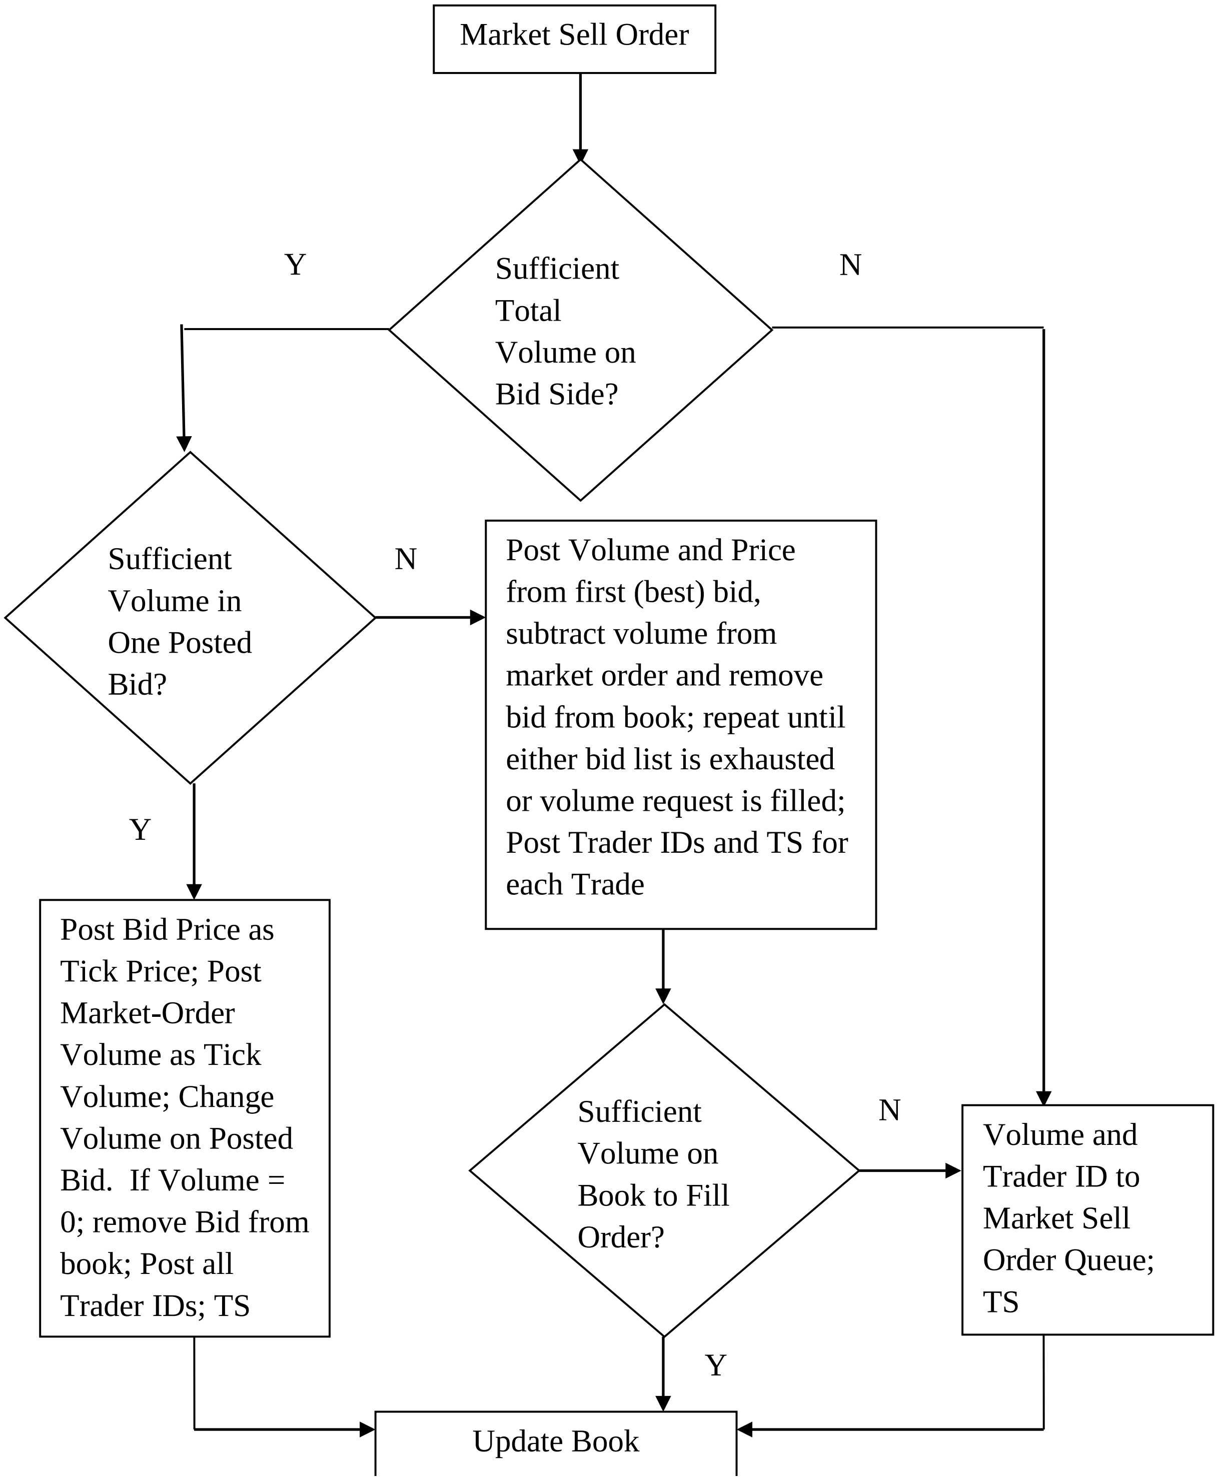
Market Sell Order
Y
Sufficient Total Volume on Bid Side?
N
Sufficient Volume in One Posted Bid?
N
Post Volume and Price from first (best) bid, subtract volume from market order and remove bid from book; repeat until either bid list is exhausted or volume request is filled; Post Trader IDs and TS for each Trade
Y
Post Bid Price as Tick Price; Post Market-Order Volume as Tick Volume; Change Volume on Posted Bid. If Volume = 0; remove Bid from book; Post all Trader IDs; TS
Sufficient Volume on Book to Fill Order?
N
Volume and Trader ID to Market Sell Order Queue; TS
Y
Update Book

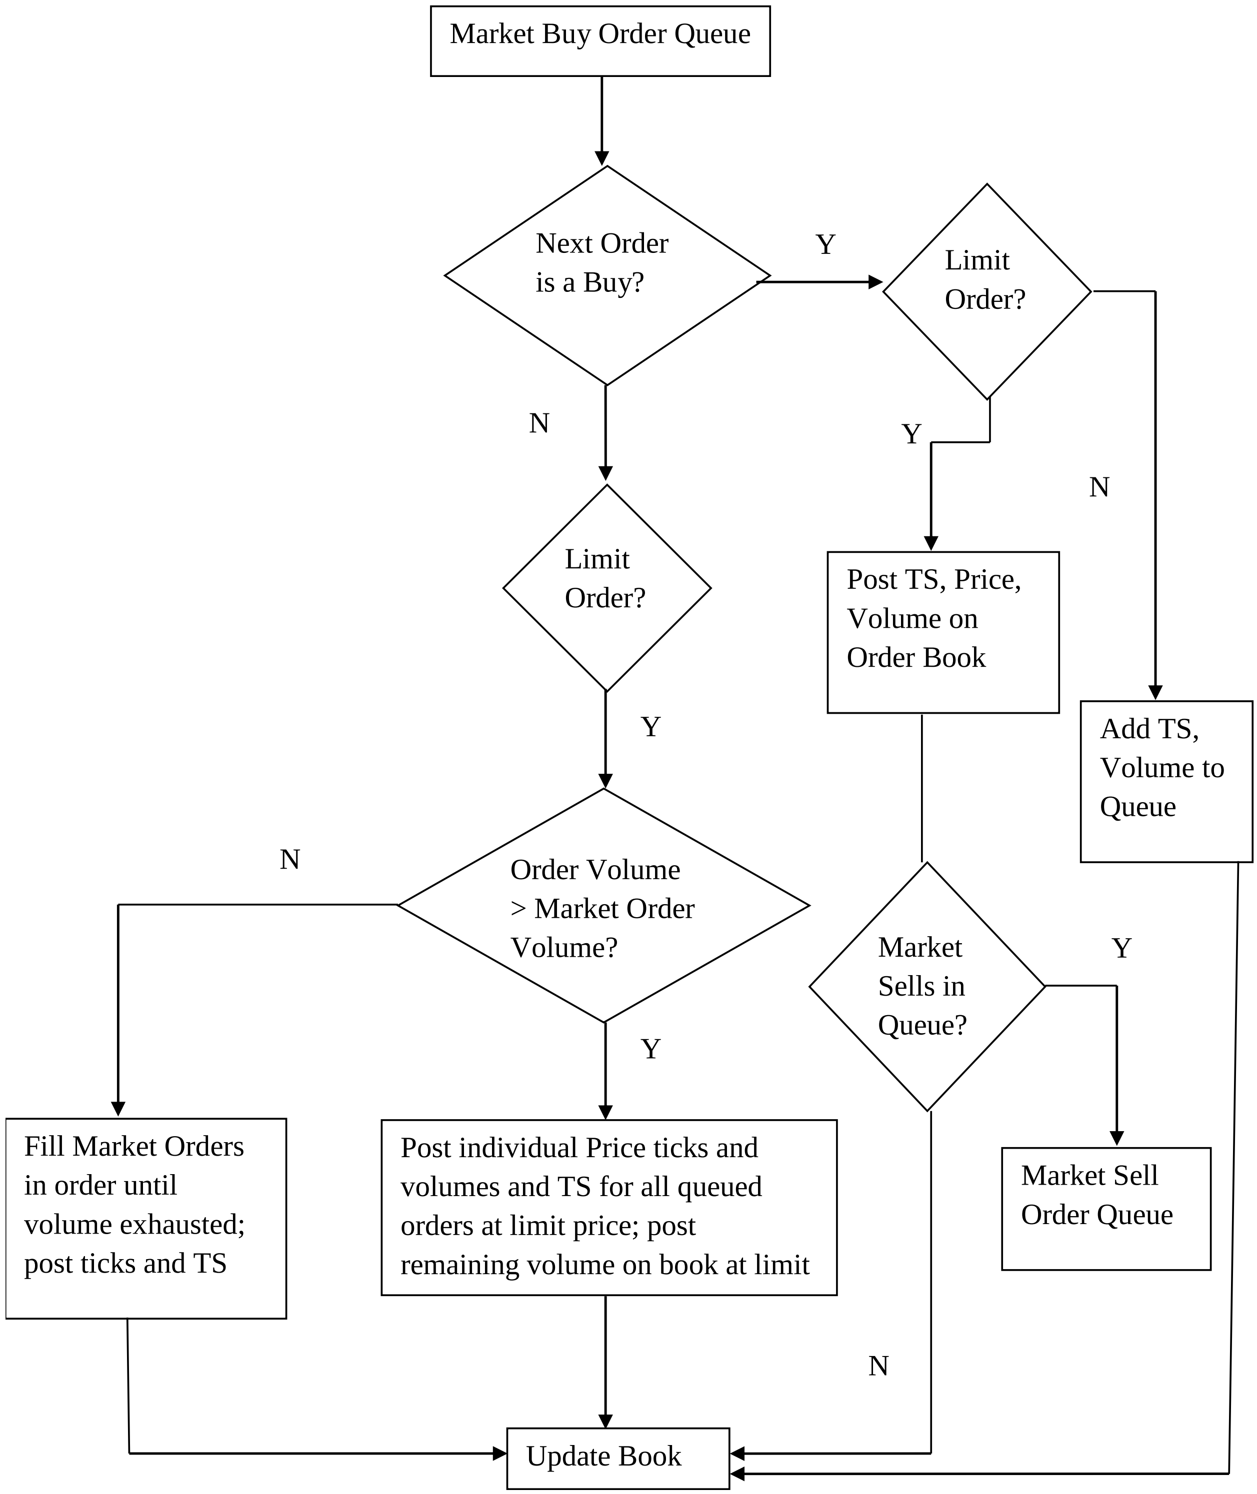

Market Buy Order Queue
Next Order is a Buy?
Y
Limit Order?
N
Y
N
Limit Order?
Post TS, Price, Volume on Order Book
Add TS, Volume to Queue
Y
N
Order Volume > Market Order Volume?
Market Sells in Queue?
Y
Y
Fill Market Orders in order until volume exhausted; post ticks and TS
Post individual Price ticks and volumes and TS for all queued orders at limit price; post remaining volume on book at limit
Market Sell Order Queue
N
Update Book

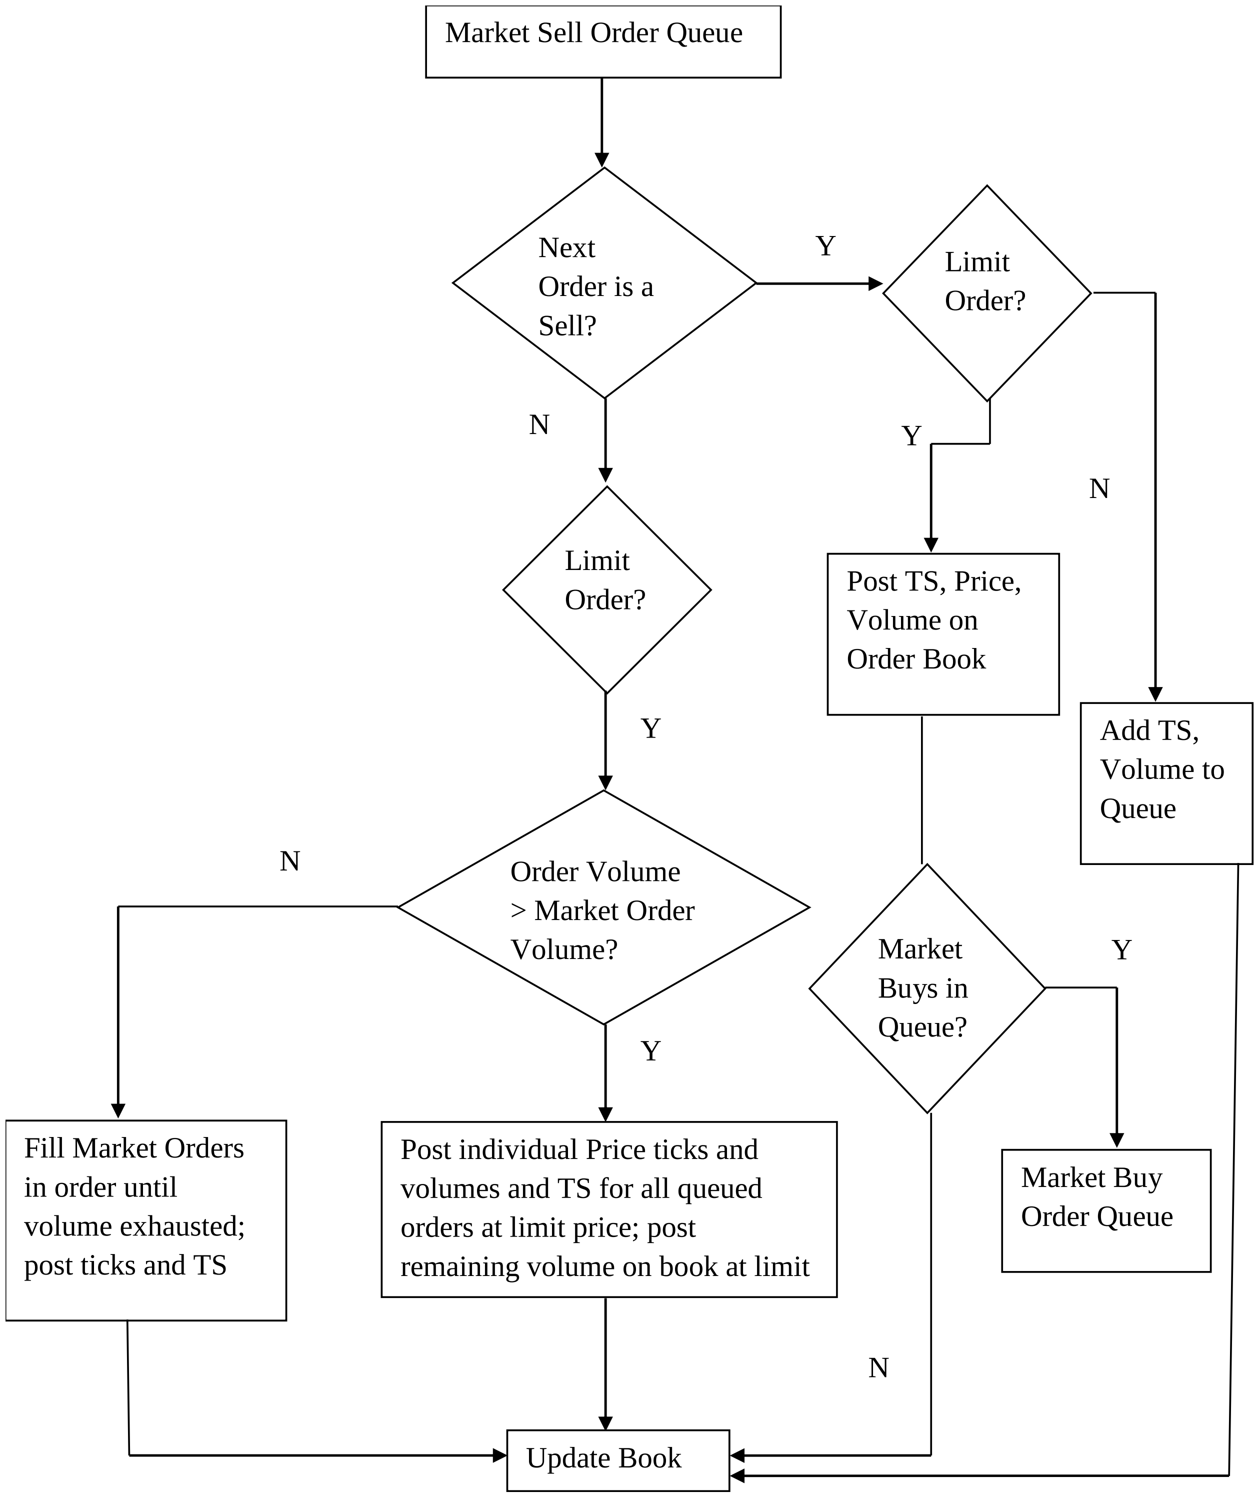

Market Sell Order Queue
Next Order is a Sell?
Y
Limit Order?
N
Y
N
Limit Order?
Post TS, Price, Volume on Order Book
Y
Add TS, Volume to Queue
N
Order Volume > Market Order Volume?
Market Buys in Queue?
Y
Y
Fill Market Orders in order until volume exhausted; post ticks and TS
Post individual Price ticks and volumes and TS for all queued orders at limit price; post remaining volume on book at limit
Market Buy Order Queue
N
Update Book

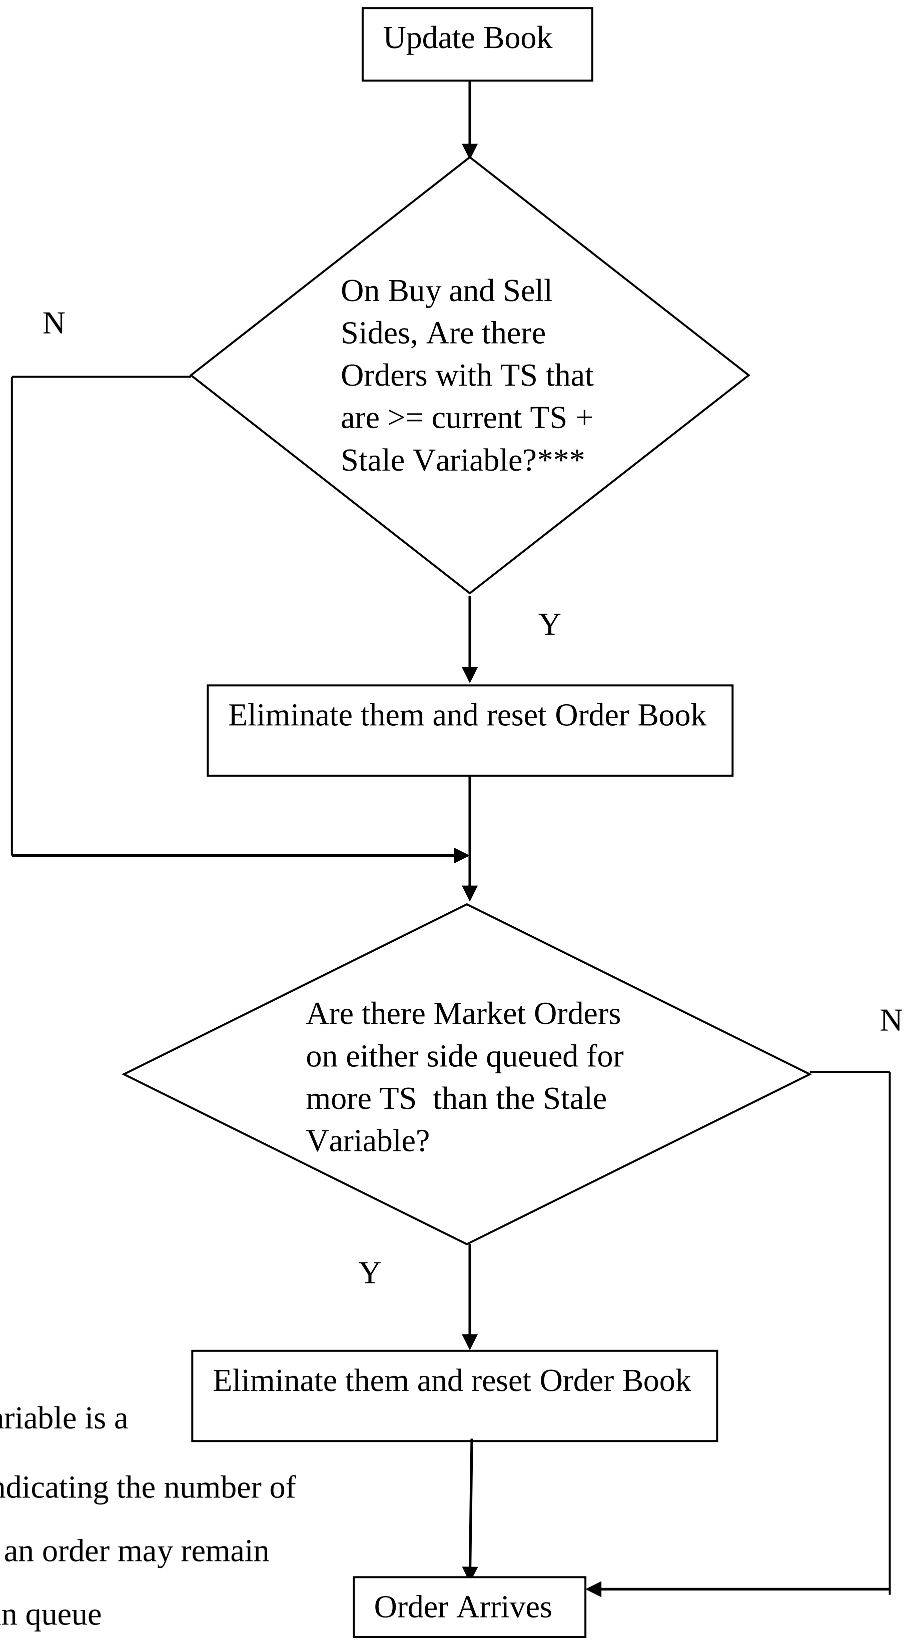

*** Stale variable is a user input indicating the number of time (steps) an order may remain on book or in queue

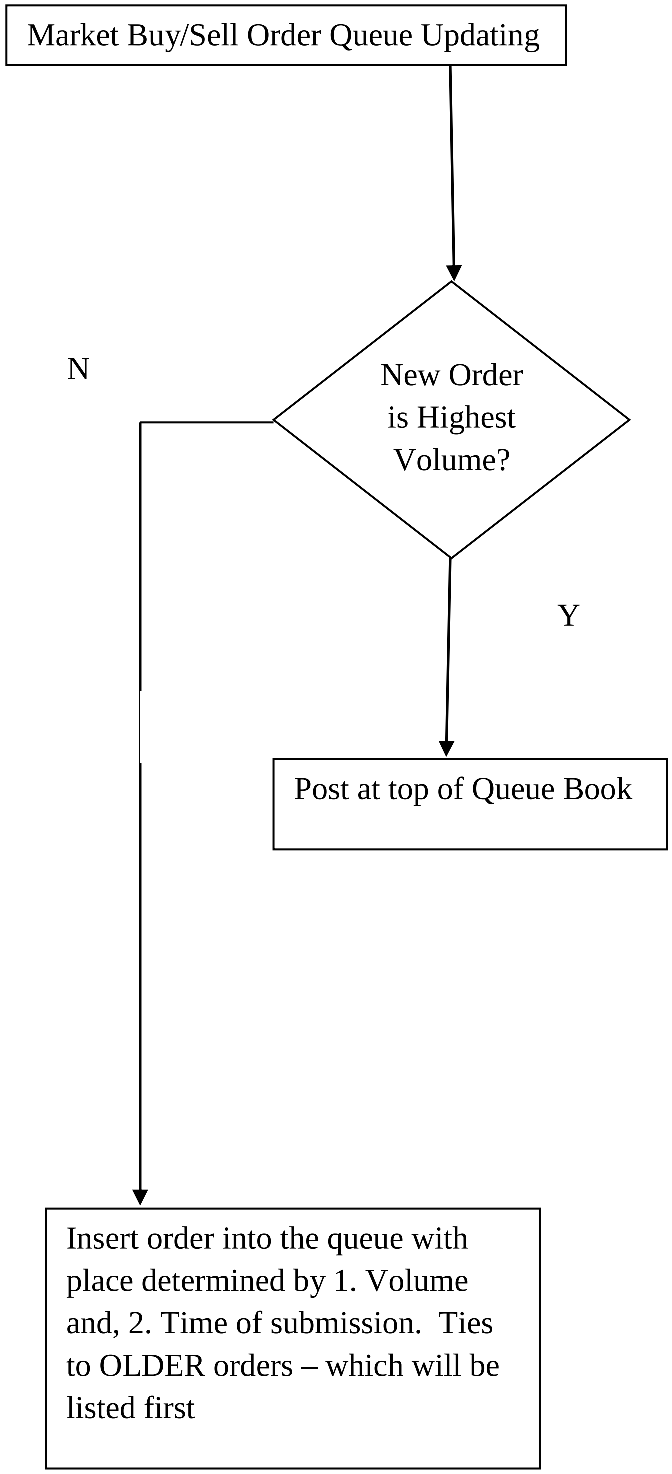
Market Buy/Sell Order Queue Updating
N
New Order
is Highest
Volume?
Y
Post at top of Queue Book
Insert order into the queue with
place determined by 1. Volume
and, 2. Time of submission. Ties
to OLDER orders – which will be
listed first

APPENDIX 2

Simulated Market Maker Decision Process

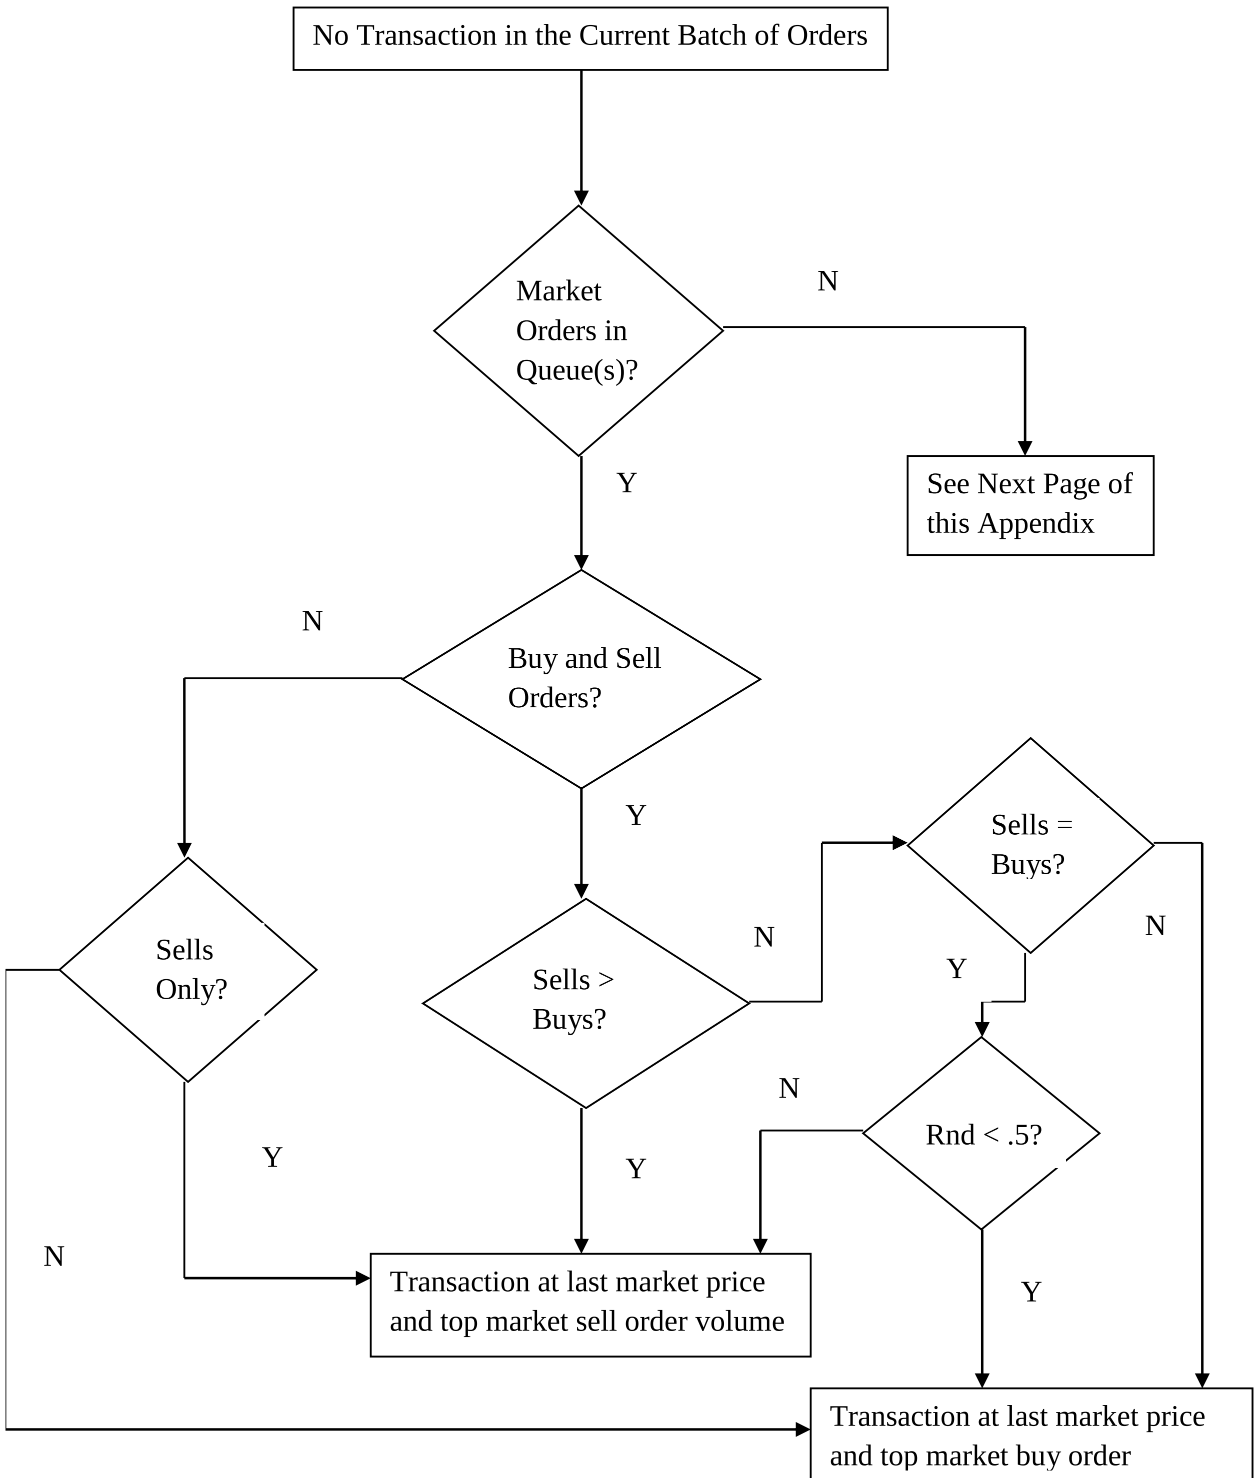

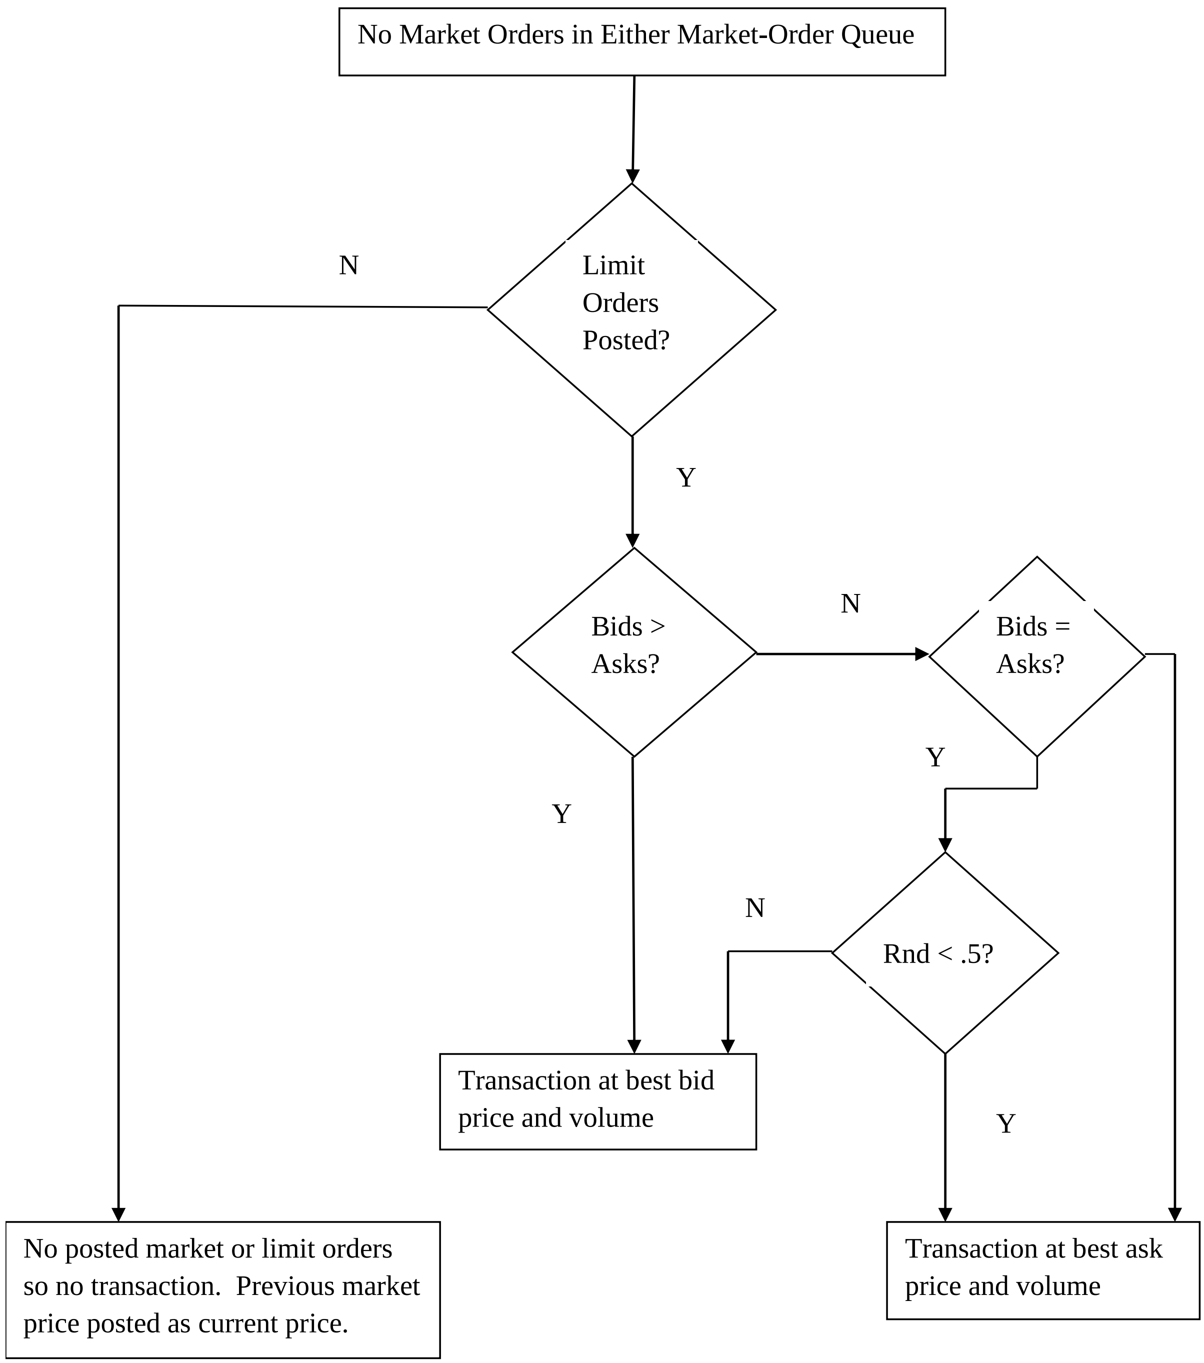
No Market Orders in Either Market-Order Queue
Limit Orders Posted?
N
Y
Bids > Asks?
N
Bids = Asks?
Y
Y
N
N
Rnd < .5?
Transaction at best bid price and volume
Y
No posted market or limit orders so no transaction. Previous market price posted as current price.
Transaction at best ask price and volume

REFERENCES

Anufriev, M. and V. Panchenko (2009): Asset Prices, Traders' Behavior, and Market Design, *Journal of Economic Dynamics and Control*, 33, 1073-1090.

Bachelier, L (1900): The Theory of Speculation, PhD Thesis, l'EcoleNormaleSuperieure. Available in English from:http://www.radio.goldseek.com/bachelier-thesis-theory-of-speculation-en.pdf

Botazzi, G., G. Dosi, and I. Rebesco (2005): Institutional Architectures and Behavioral Ecologies in the Dynamics of Financial Markets, *Journal of Mathematical Economics*, 41, 197-228.

Brock, W. A. and C. Hommes, (1998): Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model, *Journal of Economic Dynamics and Control*, 22, 8-9, 1235-1274.

Brock, W. A., C. Hommes, and F.O.O. Waggener (2005): Evolutionary Dynamics in Markets with Many Trader Types, *Journal of Mathematical Economics*, 41, 7-42.

Chen, S. H. (2012):Varieties of agents in agent-based computational economics: A historical and an interdisciplinary perspective, *Journal of Economic Dynamics and Control*, 36,1, 1-25.

Chiarella, C. and X-Z He (2003**):** Heterogeneous Beliefs, Risk, and Learning in a Simple Asset-Pricing Model with a Market Maker, Macroeconomic Dynamics, 7, 503-536.

Chiarella, C., X-Z He, and C. Hommes (2007): A Dynamic Analysis of Moving Average Rules, Journal of Economic Dynamics and Control, 30, 9-10, 1729-1753.

Chiarella, C., G. Iori, and J. Perelo (2009): The Impact of Heterogeneous Trading Rules on the Limit Order Book and Order Flows, Journal of Economic Dynamics and Control, 33, 3, 525-537.

Chiarella, C., X-Z He, and R.C.J. Zwinkels, (2014): Heterogeneous Expectations in Asset Pricing: Empirical Evidence from the S&P 500, Journal of Economic Behavior and Organization, 105, 1-16.

de Jong, E., W. F. C. Verschoor, and R.C.J. Zwikels (2010): Heterogeneity of Agents and Exchange Rate Dynamics, Journal of International Money and Finance, 29, 8, 1652-1659.

Fama, E. (1976): *Foundations of Finance*. Basic Books.

Fama, E. and K. French (1993): Common Risk Factors in Returns on Stocks and Bonds, *Journal of Financial Economics*, 33, 1, 3-56.

Fournier, M. and K. Jacobs (2016): Inventory, Risk, Market-Maker Wealth and the Variance Risk Premium: Theory and Evidence, working paper available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2334842.

He, X-Z and M. Zheng, Dynamics of Moving Average Rules in a Continuous-Time Financial Market Model, Journal of Economic Behavior and Organization, 76, 615-634.

LeBaron, B. (2002): Building the Santa Fe Artificial Stock Market

Lintner, J. (1965): The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, Review of Economics and Statistics, 47,1, 13-37.

Lux, T. and S. Alfaro (2016): Financial Power Laws: Empirical Evidence, Models, and Mechanisms, *Chaos, Solutions, and Fractals*, 88, 3-16.

Osborne, M.F.M., Brownian Motion in the Stock Market, *Operations Research* 7 (March-April), 145-173.

Raberto, M. and S. Cincotti, (2005): Modeling and Simulation of a Double Auction Financial Market, *Physica A*, 355, 34-45.

Rosu, I. (2009): A Dynamic Model of the Limit Order Book, *Review of Financial Studies*, 22, 4601-4641.

Sharpe, W.F. (1964): *Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk*, Journal of Finance, 19, 3, 425-442.

Shen, P. and R. M. Starr, Market-maker's Supply and Pricing of Financial Market Liquidity, Economics Letters, 76, 53-58.

Wang, L-X, (2005): Dynamical Models of Stock Prices Based on Technical Trading Rules, *IEEE Transactions on Fuzzy Systems*, 23, 4, 787-801.