Optimal Trading with Limit Orders on a Dynamic Limit Order Book

Craig W. Holden
Indiana University

July 2014

Abstract

I determine the optimal trading strategy for an institutional trader who wants to purchase a large number of shares over a fixed time horizon. First, I consider when a multi-period trader can submit limit orders as well as market orders. I develop a simple binomial model where limit orders either execute or not and solve it analytically. I find that the optimal sequence of limit orders involves dynamic aggressiveness. That is, if a given limit order executes (or not), then the next limit order optimally has a slightly less (more) aggressive price. I find that this trading strategy frequently beats (or at least ties) the benchmark trading strategies from the existing literature. Second, I develop an elaborate model of a multi-period trading algorithm that allows multiple orders over multiple periods to be submitted to a dynamic limit order book exchange. I calibrate the simulation to real-world summary statistics based on order data. For each trading problem, I use genetic algorithm optimization to numerically determine the optimal trading algorithm over a large number of simulated trials. I find that the optimal trading algorithm typically involves dynamic aggressiveness. Further, I find that if the fund manager is opportunistic, then the optimal algorithm involves only limit orders with low price aggressiveness. Conversely if the fund manager is committed, then limit orders should be followed by market orders at the end. I find that if the fund manager is informed and not using effective spread to measure the cost of trading, then market orders should be front-loaded in time. Conversely, if effective spread is used or the fund manager is uninformed, then less aggressive orders should be spread evenly over time. I find that these trading algorithms frequently beat (or at least tie) the benchmark trading algorithms from the existing literature.

Keywords: Optimal execution, dynamic limit order book, liquidity, order flow.
JEL classification: G14, D44.

Addresses for correspondence: Craig Holden, Kelley School of Business, Indiana University, 1309 E. Tenth St., Bloomington, IN 47405-1701. Email: cholden@indiana.edu

Acknowledgements: I thank Eitan Goldman, Noah Stoffman, Charles Trzcinka, Avi Wohl, Scott Yonker, and seminar participants at Indiana University and Tel Aviv University. I am solely responsible for any errors.

Consider an institutional trader who wishes to buy 100,000 shares of a given stock over a trading day on a limit order book exchange. What orders should the trader submit? There are many choices. First is the choice of order type: market order versus limit order.[1] A market order will get the shares requested, but at a higher cost. A limit order will cost less if it executes, but it may not execute. Second, if a limit order is used, then there is a choice of price. A higher limit buy price will be more likely to execute, but a lower limit buy price is a better price. Third is the choice of order size. Should the trader submit a smaller number of big orders or a larger number of small orders? Fourth is the choice of dynamic strategies. If new information arrives, should the unexecuted portion of a limit order be cancelled and resubmitted at an updated price? Since market orders pay the spread, should one strategically wait for moments when the spread is relatively small? And there are many more dynamic strategies that one could imagine.

This paper analyzes the optimal trading strategy of a buy-side trader who wants to purchase a large number of shares over a fixed time horizon. Specifically, I consider two research questions. First, when a multi-period trader can submit a limit or market order, is it possible to beat benchmark trading strategies from the existing literature? Second, when a multi-period trading algorithm can depend on a rich set of state variables from a dynamic limit order book model calibrated to real-world parameters, is it possible to beat trading algorithms from the existing literature?

To answer these questions, I develop two models. First, I develop a simple binomial representation of trading in partial equilibrium (i.e., where an individual optimizes while taking the rest of the world as given). In this model trading is done with unit-sized limit orders and market orders. The model is binomial in the sense that orders either fully execute or don't execute. Since

[1] A market order is a request to buy or sell a specific quantity of shares at available market price(s). A limit order is a request to buy or sell at a specified price as many shares as possible up to desired quantity.

all orders are unit-sized, there is no case of partial execution. I find the *analytic* solution for the optimal trading strategy for both single-period and multiple-period problems. I find that the optimal multi-period trading strategy involves what I call "dynamic aggressiveness." Dynamic aggressiveness means that the optimal sequence of limit orders involves small changes in price aggressiveness[2] from node-to-node over a binomial tree. That is, if a given limit order executes (or doesn't), then the next limit order optimally has a slightly less (more) aggressive price. This trading strategy frequently beats (or at least ties) the benchmark trading strategies from the existing literature.

Secondly, I develop an elaborate model of a multi-period trader who submits orders to a dynamic limit order book exchange in partial equilibrium. The buy-side trader is trying to satisfy an overall trading request by a fund manager. The model is wide-open. Limit or market orders of any size can be submitted. Prices are on a discrete penny grid. Unexecuted limit orders can be cancelled at any time. The buy-side trader can select from a huge range of trading algorithms. Trading algorithms can depend on a rich variety of state variables such as the current bid, ask, midpoint, bid depth, ask depth, order arrival from other traders, order type from other trades, order size from other traders, own quantity executed so far, number of periods remaining, lagged values of any variable, etc. I calibrate the model to real-world summary statistics for order arrival, order size, and intraday price volatility. For each trading problem, I use genetic algorithm optimization to *numerically* determine the optimal trading algorithm over a large number of simulated trials.

I find that the optimal trading algorithm typically involves dynamic aggressiveness. I find that if the fund manager is opportunistic (e.g., suffers no disutility from failure to obtain the requested quantity), then the optimal algorithm should use only limit orders with low price

[2] For limit buy (sell) orders, a higher (lower) price is more aggressive.

aggressiveness, because these trades will earn the spread. Conversely, if the fund manager is committed (e.g., strongly wants the requested quantity and so would suffer large disutility from failure to obtain it), then the optimal algorithm uses limit orders followed by market orders at the end, because the market orders will guarantee purchasing the requested amount. I find that if the fund manager is informed and performance is not measured using effective spread, then the optimal algorithm uses market orders that are front-loaded in time so as to trade before prices move in the predicted direction. Conversely, if performance is measured using effective spread or the fund manager is uninformed, then the optimal algorithm uses less aggressive orders spread evenly over time. The reason is that under effective spread the benchmark is the contemporaneous quote midpoint, which moves up or down with the price, so there is no failure penalty to trading later in the day. Also, if the fund manager is uninformed, then there isn't any predicted direction of prices to avoid. I find that the optimal trading algorithms frequently beat (or at least tie) the benchmark trading strategies from the existing literature.

This paper sits at the intersection of two streams of literature. One stream of literature analyzes the optimal execution problem of a large, multi-period trader in partial equilibrium. Bertsimas and Lo (1998), Harris (1998), Almgren and Chriss (1999, 2001), Vayanos (2001), Huberman and Stanzl (2005), Obizhaeva and Wang (2012), and others analyze a multi-period trader whose trades have temporary and/or permanent price impact in partial equilibrium.[3] These models have contributed to our knowledge of optimal multi-period trading strategies, but the need for tractability has previously limited consideration to market order strategies only. Both of my models allow both limit and market orders.

[3] Obizhaeva and Wang (2012) develop an interesting variation in which the multi-period trader's market orders deplete a block-uniform limit order book on continuous prices.

Another stream of literature theoretically analyzes limit order book exchanges. Dynamic limit order book models have been developed by Parlour (1998), Foucault (1999), Parlour and Seppi (2003), Foucault, Kadan, and Kandel (2005), Goettler, Parlour, and Rajan (2005, 2009), Rosu (2009), Colliard and Foucault (2012), and Buti and Rindi (2013). These models have contributed useful characterizations of limit order book markets, but their significant complexity has required the simplification that each endogenous agent be limited to a single period for order submission. Both of my models allow an endogenous trader to submit multiple orders over multiple periods.

The paper is organized as follows. Section 1 develops a simple binomial model of multi-period trading and solves it analytically. Section 2 develops an elaborate model of a multi-period trading via a dynamic limit order book, calibrates it, and numerically solves for the optimal trading algorithm. Section 3 concludes. An appendix contains all proofs.

1. A Simple Binominal Model

1.1 The Single-Period Version

I develop a simple binomial model of a pure, open limit order market. This type of market structure is used by NYSE-ARCA, BATS, Direct Edge, the Tokyo Stock Exchange, Euronext, and many other exchanges around the world. This simple model provides insight and intuition into optimal trading strategy that carries over to the rich simulation model in the next section of the paper.

I begin with a single-period version. For simplicity, the initial limit order book for a given stock is empty. A single seller wishes to trade with a single buyer. Both traders need to set a limit order price. By convention, the seller is exogenous. The seller submits a limit sell order at a limit sell price s for one unit of the asset. The limit sell price s is drawn from an uniform distribution over the continuous interval $[v-d, v+d]$, where v is the public value of the stock and d specifies the amount of dispersion of potential prices. The buyer is endogenous. The buyer submits a limit buy order at a limit buy price b for one unit of the asset. He chooses the limit buy price b from the same the continuous interval $[v-d, v+d]$. Define limit buy price aggressiveness a, as the excess of the limit price above the public value, $a = b - v$.

Both orders are submitted simultaneously. Both orders execute in full when $s \le b$ and both fail to execute when $s > b$. It immediately follows that the limit buy's probability of execution p is the probability that $s \le b$, which is

$$p = \frac{b-(v-d)}{(v+d)-(v-d)} = \frac{b-v+d}{2d}. \qquad (1)$$

This way of modeling a limit buy captures the key property that a higher buy price b increases the probability of execution p. At the extreme, a buy price of $v+d$ yields a 100% probability of execution. I interpret this case as a market buy order.[4] For simplicity, I adopt the convention that the two orders execute at the buy price b.[5]

The buyer wishes to minimize his expected disutility $E[D]$ at the end of the period. If the limit buy executes, then he receives disutility from the cost of trading $b-v$, the excess price paid relative to the public value v.[6] If the limit buy fails to execute, then he receives a failure penalty

[4] Equivalently, this can be interpreted as a marketable limit buy. The simple binomial model does not distinguish between a market buy and a marketable limit buy.
[5] It is easy to show that the results of the model are qualitatively similar for the alternative convention that they execute half the time at the buy price and half the time at the sell price.
[6] This measure of the cost of trading is very similar to the effective half spread, which for buy trades is defined as the trade price minus the quote midpoint. In this simple model, the cost of trading is equal to price aggressiveness, $a = b - v$. In the rich simulation model in the next section, we will measure the cost of trading with four different metrics.

k, which represents the buyer's disutility for failing to obtain the desired amount. The buyer's decision problem is to choose the limit buy price b so as to minimize his expected disutility

$$\underset{b}{Min}\, E[D] = p(b-v) + (1-p)k \,. \tag{2}$$

A higher b increases p which reduces the chance of receiving the failure penalty k for failing to execute. The tradeoff is that a higher b increases the cost of trading $b-v$ when execution happens. The proposition below gives the analytic, single-period solution.

Proposition 1. If $k < 3d$ *, then the optimal limit price aggressiveness is*

$$a = \frac{1}{2}(k-d) < d \,, \tag{3}$$

the corresponding probability of execution is

$$p = \frac{k+d}{4d} < 1 \,, \tag{4}$$

and the corresponding expected disutility is

$$E[D] = \frac{6dk - d^2 - k^2}{8d} < d \,. \tag{5}$$

If $k \geq 3d$ *, then the optimal limit price aggressiveness is* $a = d$ *, the corresponding probability of execution is* $p = 1$*, and the corresponding expected disutility is* $E[D] = d$*.*

The key driver of the optimal trading strategy is the failure penalty k . If the failure penalty is low enough $(k < 3d)$, then the optimal order submission is a limit buy with an interior optimum price aggressiveness $a < d$. The price aggressiveness is increasing in the failure penalty k and decreasing in the dispersion d. Conversely, if the failure penalty is high enough $(k \geq 3d)$, then the optimal order submission is a market buy at the corner solution price aggressiveness $a = d$.

1.2 The Multi-Period Version

Now I extend the simple binomial model to a multi-period version. There are T order submission times $t = 1, 2, \ldots, T$ and a terminal time $T+1$.[7] Let v_t be the public value of the stock at time *t*. This public value is conditional on all public information at time *t*, including the history of all prior trades and quotes in the stock. Let next time's public value be given by $v_{t+1} = v_t + \varepsilon_{t+1}$, where ε_{t+1} is next time's innovation in the pubic value.

At any node on a binomial tree, a single exogenous seller wishes to trade immediately. At time t, the seller submits a limit sell at a sell price s_t for one unit that expires after one time period. The limit sell price s_t is drawn from an uniform distribution over the continuous interval $[v_t - d, v_t + d]$.

A single endogenous long-lived buyer wishes to purchase N units by the terminal time $T+1$. Since there is only one unit available for sale at a given time, this will require the execution of N limit or market buy orders of one unit each. We constrain $N \leq T$ so that it is feasible in some state of nature to purchase the desired quantity. We also impose the economically sensible non-negativity constraints $d > 0,\ N > 0,$ and $v_t - d > 0$ at all times.

Consider a buyer who at time *t* has already succeeded in purchasing n units. At the (t, n) node of a binomial tree, the buyer cancels any unexecuted limit buy from the previous time[8] and submits a new limit buy order at a buy price $b_{t,n}$ for one unit. He chooses $b_{t,n}$ from the same

[7] In principle, *T* could be arbitrarily large. There could be so many periods that one period could represent one second or even one millisecond.

[8] In the simple binomial model, canceling any unexecuted limit buy always leaves the buyer either better off or indifferent compared to having two limit buys outstanding at the same time. To see this, recall that only one unit is available for sale in any period and thus, it is not possible for two buy orders to execute in the same period. Further, if the prior buy price is higher than the current period optimal buy price, then the prior buy is too high and the buyer is better off by canceling it. If the prior buy price is lower than the current optimal buy price, then the prior buy price is dominated and the buyer is indifferent to its existence. So in the latter case, we adopt the tie-breaking convention that he cancels it.

interval $\left[v_t - d, v_t + d\right]$ as the contemporaneous seller. As before, the limit buy price aggressiveness $a_{t,n}$ is defined as the excess of the limit price above the public value, $a_{t,n} = b_{t,n} - v_t$.

Both orders are submitted simultaneously. Both orders execute in full when $s_t \le b_{t,n}$ and both fail to execute otherwise. The limit buy's probability of execution $p_{t,n}$ is the probability that $s_t \le b_{t,n}$, which is given by

$$p_{t,n} = \frac{b_{t,n} - \left(v_t - d\right)}{\left(v_t + d\right) - \left(v_t - d\right)} = \frac{b_{t,n} - v_t + d}{2d}. \tag{6}$$

The long-lived buyer's disutility function at the terminal $\left(T+1, n\right)$ node is

$$D_{T+1,n} = TCT_{T+1} + k\left(N - n\right) \qquad \text{for } n = 0,1,\ldots,N, \tag{7}$$

where TCT_{T+1} is the total cost of trading prior to time $T+1$ (which includes the final trading time *T*) and the second term is the trader's degree of disutility due to the quantity underfill $N - n$.

Let $J_{t,n}$ be the long-lived buyer's derived disutility function at the $\left(t, n\right)$ node. By dynamic programming, the decision problem at the $\left(t, n\right)$ node is given recursively by

$$J_{t,n} = \underset{b_{t,n}}{Min}\, E_{t,n}\left[J_{t+1,\tilde{n}}\right]. \tag{8}$$

where $\tilde{n}$ is the random number of units purchased through the next time $t+1$ and $J_{T+1,n} = D_{T+1,n}$ for all values of *n*. At any given $\left(t, n\right)$ node, if $n < N$, then an order will be submitted that will either execute or not. However, if $n = N$, then no order will be submitted since the full desired quantity has already been purchased. Thus, derived disutility $J_{t,n}$ at the $\left(t, n\right)$ node can be written as

$$J_{t,n} = \underset{b_{t,n}}{Min} E_{t,n}\left[J_{t+1,\tilde{n}} \right] = \begin{cases} \underset{b_{t,n}}{Min}\ p_{t,n} J_{t+1,n+1} + \left(1 - p_{t,n}\right) J_{t+1,n} & \text{if } n < N \\ TCT_t & \text{if } n = N. \end{cases} \tag{9}$$

In the case that $n = N$, the absence of further orders means that the current derived utility equals the terminal disutility, which also equals the total cost of trading that has already been realized prior to the current date TCT_t and no failure penalty due to quantity underfill.

In the case that $n < N$, we can write the derived disutility as the expected value of terminal disutility under the optimal trading strategy as determined recursively backwards by dynamic programming. In the event that the time *t* order executes, then the updated expected value of terminal disutility is

$$J_{t+1,n+1} = TCT_t + b_{t,n} - v_t + f_{t+1,n+1} + k\left(N - (n+1) - h_{t+1,n+1}\right), \tag{10}$$

where TCT_t is the total cost of trading prior to time *t*, $b_{t,n} - v_t$ is the increment to the total cost of trading caused by the time *t* order executing, $f_{t+1,n+1}$ is defined as the expected future addition to the total cost of trading starting from the $\left(t+1, n+1\right)$ node, and $h_{t+1,n+1}$ is defined as the expected number of units to be purchased in future trades starting from the $\left(t+1, n+1\right)$ node. Both $f_{t+1,n+1}$ and $h_{t+1,n+1}$ are the forecasted outcome of future trades under the optimal trading strategy.

In the event that the time *t* order doesn't execute, then the updated expected value of terminal disutility is

$$J_{t+1,n} = TCT_t + f_{t+1,n} + k\left(N - n - h_{t+1,n}\right), \tag{11}$$

where $f_{t+1,n}$ and $h_{t+1,n}$ are defined analogously starting from the $\left(t+1, n\right)$ node.

Again there is a trade-off. A higher buy price $b_{t,n}$ increases the probability of execution $p_{t,n}$ which reduces the chance of receiving the failure penalty (disutility) for a quantity underfill, but it increases the increment to the total cost of trading $b_{t,n} - v_t$ when execution does happen. The proposition below gives the analytic, multi-period solution.

Proposition 2. The solution is specified by binomial trees for $a,\ p,\ f,$ and *h. At the* (t,n) *nodes of these trees where the trader is not done* $(n < N)$*, the optimal limit price aggressiveness is*

$$a_{t,n} = Min\left[d, \frac{1}{2}\left(f_{t+1,n} - f_{t+1,n+1} + k\left(1 + h_{t+1,n+1} - h_{t+1,n}\right) - d \right) \right], \tag{12}$$

the probability of execution at the optimal limit buy price is

$$p_{t,n} = \frac{a_{t,n} + d}{2d}, \tag{13}$$

the expected total cost of trading on future trades is

$$f_{t,n} = p_{t,n}\left(b_{t,n} - v_t + f_{t+1,n+1}\right) + \left(1 - p_{t,n}\right) f_{t+1,n}, \tag{14}$$

and the expected number of units to be purchased in future trades is

$$h_{t,n} = p_{t,n}\left(1 + h_{t+1,n+1}\right) + \left(1 - p_{t,n}\right) h_{t+1,n}. \tag{15}$$

At the (t,n) *nodes of these trees where the trader is done* $(n \geq N)$*, no further orders will be submitted and so* $f_{t,n} = 0,\ h_{t,n} = 0,\ p_{t,n} = 0,$ *and* $a_{t,n}$ *is undefined. For all terminal nodes,* $f_{T+1,n} = 0,$ and $h_{T+1,n} = 0.$ *The* a *and* p *trees do not include the terminal date* $T+1$*. Further, a binomial tree for the derived disutility minus the total cost of trading prior to time t* $\left(J_{t,n} - TCT_t\right)$ *is given by equations (9), (10), and (11). On the initial date, the total cost of trading before the first trade* $TCT_1 = 0$ *and so the first node (1,0) of this binomial tree yields the ex-ante derived disutility* $J_{1,0}$.

The binomial trees are calculated backwards from the last date to the first. Start at the date $T+1$ nodes for the f and h binomial trees (which are equal to zero), next calculate the date T nodes for the a and p binomial trees using (12) and (13), then calculate the date T nodes for the f and h trees using (14) and (15), …, keep calculating the a and p trees first and then the f and h trees second on each date all the way back to date 1. Finally, the $J-TCT$ binomial tree can be calculated from the a, p, f, and h trees using (9), (10), and (11) starting at date T and working back to date 1.

1.3 The Character of the Optimal Trading Strategy

A numerical example will illustrate the character of the optimal trading strategy. Suppose that the long-live buyer wishes to purchase 4 units $(N=4)$ has 8 time periods $(T=8)$ to do so. Further suppose that the remaining parameters are $v_0=\$30.00,$ $d=\$0.10,$ and $k=0.32.$ Given these parameters, the failure penalty is relatively high $(k>3d)$.

Figure 1 shows a binomial tree for the optimal limit price aggressiveness a and a binomial tree for the probability of execution p. Starting at the (1,0) node of the upper tree, the optimal limit price aggressiveness on date 1 is $a_{1,0}=\$0.00,$ meaning the date 1 limit buy price is set equal to the date 1 true value $b_{1,0}=v_1$. Looking at the corresponding (1,0) node of the lower tree, this order has a probability of execution $p_{1,0}=50\%$. If this order *fails* to execute, then you go to the *up*-node (2,0) in both trees. At this node, you cancel the time 1 limit order and submit an optimal time 2 limit order with a slightly more aggressive price $a_{2,0}=\$0.01$ (meaning a limit buy price $b_{2,0}=v_2+.01)$, which corresponds to a slightly higher probability of execution $p_{2,0}=56\%$. However, if the time 1 order *succeeds* in executing, then you go to the *down*-node (2,1) in both trees. At this node, you submit an optimal time 2 limit order with a slightly less aggressive price

$a_{2,1} = -\$0.01$ (meaning a limit buy price $b_{2,1} = v_2 - .01$), which corresponds to a slightly lower probability of execution $p_{2,1} = 44\%$.

I call this pattern of slightly higher price aggression on up-steps and slightly lower price aggression on down-steps "dynamic aggressiveness." That is, from any node if the current limit order fails to execute, then the next step at the up-node will optimally be at a slightly more aggressive price with a slightly higher probability of execution. Intuitively, the failure to execute makes the remaining problem *more* difficult (you have one less period to purchase the same number of units as before) and so you act *more* aggressively. Conversely, if the current limit order succeeds in executing, then the next step at the down-node will optimally be at a slightly less aggressive price with a slightly lower probability of execution. Intuitively, the success at execution makes the remaining problem *less* difficult (you have one less unit to purchase in the remaining time) and so you act *less* aggressively.

Dynamic aggressiveness characterizes nearly all of the 2^T possible paths that you might take over the binomial tree. For example, suppose your first four steps happen to be up, down, down, and up. Then, your price aggressiveness and probability of execution will slightly increase, decrease, decrease, and increase.

More specifically, dynamic aggressiveness characterizes any path over the entire diamond-shaped zone with the yellow shading (light grey shading with solid borders in black and white) on the binomial trees. I call this area the Limit Order zone, because it turns out that it is always optimal to submit non-marketable limit orders $(p_{t,n} < 1)$ in this zone (this result is proven in Proposition 4). Specifically, the Limit Order Zone contains all (t, n) nodes where the trader is not done $(n < N)$ and where there is a positive amount of slack time (i.e., more time periods left than remaining units to be purchased $T+1-t>N-n$). The following proposition formally specifies this characterization.

Proposition 3. In the Limit Order Zone, the optimal strategy exhibits dynamic aggressiveness. That is, from any (t,n) node in this zone:

- *If the order fails to execute, then you go to the up-node $(t+1,n)$ where the optimal strategy is to cancel the unexecuted order and submit a new limit buy at a more aggressive price $a_{t+1,n}>a_{t,n}$ with a higher probability of execution $p_{t+1,n}>p_{t,n}$.*
- *If the order succeeds in executing, then you go to the down-node $(t+1,n+1)$ where the optimal strategy is to submit a limit buy at a less aggressive price $a_{t+1,n+1}<a_{t,n}$ with a lower probability of execution $p_{t+1,n+1}<p_{t,n}$.*

To get a feel for dynamic aggressiveness, take a look at the lower binomial tree for the probability of execution in Figure 1. You can readily see the probability of execution at any node in the Limit Order Zone is an intermediate value between the following two connected nodes. Formally, one can make a stronger statement. Part of the proof of Proposition 3 (see the appendix) is to prove that the probability of execution at any node where the trader is not done $(n<N)$ is a weighted average of the following two connected nodes

$$p_{t,n}=w_{t,n}p_{t+1,n+1}+\left(1-w_{t,n}\right)p_{t+1,n}, \tag{16}$$

where the weight $w_{t,n}=\left(p_{t+1,n+1}+p_{t+1,n}\right)/2$. Logically, if the (t,n) node is in the middle, then one of the following connected nodes will have a higher probability of execution and the other will have a lower probability of execution.

An interesting case is when there have been four failures to execute in a row (the up-up-up-up-node). This is the (5,0) node, where there are only four time periods left and there remains four units to be purchased. In other words, there is zero slack time. The only way to guarantee execution is with the "corner solution" strategy of submitting four market buys over the next four time periods. The optimal limit price aggression for these four orders is the maximum possible $a_{5,0} = a_{6,1} = a_{7,2} = a_{8,3} = d = \$0.10$ (meaning the most aggressive possible sequence of bid prices $b_{5,0} = v_5 + d$, $b_{6,1} = v_6 + d$, $b_{7,2} = v_7 + d$, and $b_{8,3} = v_8 + d$). These market buys will execute with certainty $p_{5,0} = p_{6,1} = p_{7,2} = p_{8,3} = 100\%$. I call the upper right triangle with green shading (medium grey shading in black and white) of the binominal trees the "Market Order / Maximum Aggressiveness Zone." The name describes the optimal trading strategy in this zone (see Proposition 4 below).[9]

Another interesting case is when there have been four successes in executing in a row (the down-down-down-down-node). This is the (5,4) node, where the entire desired quantity has already been purchased. Hence, no more orders are submitted and the probability that an order will be executed is zero $p_{5,4} = p_{6,4} = p_{7,4} = p_{8,4} = 0\%$. I call the lower right triangle with red shading (dark grey shading in black and white) of the binominal trees the "Done" zone.[10]

---

[9] When the failure penalty is high $(k > 3d)$, the nodes *above* the southwest edge of the Market Order / Maximum Aggressiveness Zone are unreachable. To see this, consider the (5,0) node which is *on* the southwest edge. Here it is optimal to submit a market buy that is certain to execute $p_{5,0} = 100\%$. Thus, it is certain that the trader will go to the down-node (6,1) next and cannot reach the up-node (6,0).

[10] The nodes *below* the NorthWest edge of the Done Zone are unreachable. To see this, consider the (5,4) node which is *on* the NorthWest edge. Here no more orders are submitted and so the probability that an order will be executed is zero $p_{5,4} = 0\%$. Thus, it is certain that the trader will go to the up-node (6,4) next and cannot reach the down-node (6,5).

Now consider an example where the failure penalty is relatively low $(k < 3d)$. Reduce the failure penalty to $k = 0.28$ and keep the rest of the parameters the same: $N = 4,\ T = 8,\ v_0 = \$30.00,$ and $d = \$0.10.$ Figure 2 shows the binomial trees in this case. The main difference in this case is in the Market Order / Maximum Aggressiveness Zone. Here the optimal limit price aggression is less than the maximum possible $a_{t,n} < d$ (meaning the bid prices are less than the maximum possible $b_{t,n} < v_t + d$). These non-marketable limit orders have a probability of execution less than 100%, and specifically in this case, $p_{t,n} = 95\%.$ This limit order exhibits the Maximum Aggressiveness that is acceptable to this trader and it is greater than the aggressiveness of any order in the Limit Order Zone.[11]

Indeed, it is interesting that all of the nodes in this zone have a probability of execution which is *identical* to probability of execution for the same parameter values in single-period case. This makes sense when you consider that all of the nodes on the last trading date, (8,0), (8,1), (8,2), and (8,3), in fact only have one trading opportunity left and so logically the probability of execution $\left(p_{8,0} = p_{8,1} = p_{8,2} = p_{8,3} = 95\%\right)$ must be the identical to the single-period solution. Further, the remainder of nodes in this zone have two following connected nodes with a 95% probability of execution. Since all nodes are a weighted average of the following two connected nodes by equation (16), they must have the same 95% probability of execution.[12]

The proposition below specifies that the character of the optimal solutions that we have discussed for these particular numerical examples holds in general.

---

[11] When the failure penalty is low $(k < 3d)$, the optimal limit orders in the Market Order / Maximum Aggressiveness Zone have a probability of execution of less than 100%. Hence, all of the nodes can be reached.

[12] The situation changes when we get to the *limit order zone*, because at least one of the following connected nodes has a lower probability of execution. For example, the (7,4) node in the limit order zone is connected to the (8,4) node with a 95% probability and to the (8,5) node with a 0% probability. The weighted average of these two following connected nodes yields a 50% probability for the (7,4) node.

Proposition 4. The binomial trees for the optimal limit price aggressiveness and the probability of execution have three zones:

1. *A Limit Order Zone, where it is optimal to submit a non-marketable limit order.*
2. *A Market Order / Maximum Aggressiveness Zone. In this zone, if the failure penalty is high enough* $(k \geq 3d)$*, then it is optimal to submit a market order. Conversely, if the failure penalty is low enough* $(k < 3d)$*, then it is optimal to submit a non-marketable limit order with higher price aggressiveness than any node in the limit order zone.*
3. *A Done Zone, where the trader is done* $(n \geq N)$ *and so no further orders are submitted.*

Under dynamic aggressiveness, if the (t,n) node order fails to execute, then the remaining trading problem is more difficult and thus requires a more aggressive strategy. Pursuing this reasoning, multiple failures to execute cause the trader to cross-over into a *more aggressive zone*, namely, the Market Order / Maximum Aggressiveness Zone. Intuitively, all of the slack time is gone and so the trader must act very aggressively. Conversely, N successes in executing cause in the trader to cross-over into a *less aggressive zone*, namely, the Done Zone. Here the trader adopts the least aggressive strategy possible: stopping trading.

Next I compare the optimal strategies identified so far to two benchmark strategies from the existing literature. First, Bertsimas and Lo (1998) analyze a strategy of submitting equal-sized market orders each period. They prove that this is the optimal strategy in their framework in the special case when the underlying price innovations are zero mean independently and identically distributed random shocks (i.e., white noise). Second, Harris (1998) finds that it is optimal for a trader to start with (equal-sized) limit orders in order to attempt to get a lower price, but if this fails, then submit a market order at the end to guarantee execution.[13]

Since these two benchmark strategies are different than the dynamic aggressiveness strategy that Proposition 3 proves is optimal strategy, they must yield worse (i.e., higher) or the same ex-ante expected disutility in all cases. One tie case is the one that Harris (1998) considers, which can be viewed as a special case of the binomial model when only a single order is requested $(N=1)$ and the failure penalty is high $(k \geq 3d)$. Figure 3 shows the binomial trees in this case. Here the optimal strategy begins by submitting a limit order at a low -$0.06 aggressiveness with an 18% chance of execution. If it executes, you are done. If not, the next limit order will be more aggressive with a 20% chance of execution. The pattern continues at each node going along the upper edge of the binomial tree. From each node, you will be done if your order executes or you will submit a more aggressive next time if it fails. In the last period, you submit a market order. In this special case, the optimal strategy *exactly matches* the second benchmark from the existing literature. But in all other cases, the dynamic aggressiveness strategy beats the two benchmark strategies from the existing literature.

2. A Rich Simulation Model

2.1 Model Setup

Many buy-side institutions separate the "selection" task and the "implementation" task. For each fund (e.g., mutual fund, pension fund, endowment fund, etc.), there is a fund manager who decides what securities to buy and sell in what overall amounts at what times, which is called selection. The fund manager's trading requests are sent to a buy-side internal trader. This employee

[13] Harris and Hasbrouck (1996) also discuss the idea that a patient liquidity trader might begin by submitting a limit order, and if need be, switch to a market order at the end. Handa and Schwartz (1996) also analyze a limit order first and market order last strategy compared to a market order first strategy.

decides what order types at what order prices in what order sizes should be sent to what exchanges or broker-dealers at what times, which is called implementation.

I develop a rich simulation model from the point of view of the internal trader. This employee takes the fund manager's request as given and tries to figure out the best way to implement the request so as to minimize the fund manager's disutility function (see below).

The fund manager's request is a request to buy a specific quantity of a particular security expressed as a percentage of average daily volume.[14] For example, a request to purchase 1% of average daily volume would be an easy request, whereas a request to purchase 30% of average daily volume would be a difficult request. Multiplying (Requested Buy as a percentage of Average Daily Volume) times (Average Daily Volume measured in round lots) yields the requested number of round lots N.[15]

There are two types of agents. One agent is the endogenous, long-lived, buy-side trader. Other agents are exogenous, short-lived traders who can submit a variety of orders.[16]

The trading day is divided into 20 order submission times $t = 1, 2, \ldots, 20$ that are 20 minutes apart. Let T be the fund manager's order submission deadline to fulfill the purchase request and $T+1$ be a terminal valuation time after the deadline. For example, a fund manager's deadline of $T = 1$ (corresponding to 9:40 a.m.) would be very impatient, whereas a fund manager's deadline of $T = 20$ (corresponding to 4:00 p.m.) would be very patient.

---

[14] I only consider buy requests. Sell requests are incorporated implicitly as the mirror image of buy requests. I only consider requests to trade in a single security. Bertsimas and Lo (1998) show that portfolio requests are readily accommodated by accounting for the multivariate stochastic process across securities.

[15] One round lot is equal to 100 shares.

[16] Like the rest of the optimal execution literature, both of my models are based on partial equilibrium (that is, individual optimality). Compared to general equilibrium, the strength of this approach is its enormous generality. By calibrating the actions of other agents to real data, it avoids having to impose strong assumptions, such as all agents in the economy are rational, have full understanding of distributions, have sophisticated computational ability, etc. The inherent weakness of this approach is that it cannot address the interaction of multiple strategic agents.

The two classes of agents alternate in submitting orders. At time $t = 1$, other traders submit an order and then the buy-side trader submits an order. At time $t = 2$, other traders submit an order and then the buy-side trader submits an order. And so on until the deadline time $t = T$, when other traders submit their final order by the deadline and then the buy-side trader submits a final order by the deadline.

The fund manager's disutility function D_{T+1} at the terminal valuation time $T+1$ is

$$D_{T+1} = TCT_{T+1}^{m} + U \cdot Max\left(N - n_{T+1}, 0\right) + O \cdot Max\left(n_{T+1} - N, 0\right) + A\sigma \,, \tag{17}$$

where TCT_{T+1}^{m} is the total cost of trading at the terminal valuation time based on metric m (explained below), U is the fund manager's failure penalty (i.e, the fund manager's underfill disutility per round lot for getting too little), N is the number of round lots requested by the fund manager, n_{T+1} is the cumulative number of round lots purchased by $T+1$, O is the fund manager's overfill disutility per round lot for getting too much, A is the fund manager's risk aversion, and σ is the standard deviation of trade prices. The second term is the fund manager's degree of disutility due to a quantity underfill (when $n_{T+1} < N$), the third term is the fund manager's degree of disutility due to a quantity overfill (when $n_{T+1} > N$), and the fourth term captures the fund manager's disutility due to execution price risk. The parameters U, O, and A are required to be non-negative. In practice, we do not observe algorithms overfilling the fund manager's request. Therefore, I set O to a relatively large value such that the optimal algorithm always avoids overfilling.

I incorporate four alternative metrics that are widely-used in practice for measuring the cost of trading C_k^m due to the k^{th} trade under metric *m*:

$$Effective\ Spread: \quad C_k^{ES} = 2I_k\left(P_k - M_k\right), \tag{18}$$

$$Implementation\ Shortfall: \quad C_k^{IS} = 2I_k\left(P_k - M_{IS}\right), \tag{19}$$

$$Volume-Weighted\ Average\ Price: \quad C_k^{VWAP} = 2I_k\left(P_k - P_{VWAP}\right), \tag{20}$$

$$Closing\ Price: \quad C_k^{C} = 2I_k\left(P_k - P_C\right), \tag{21}$$

where I_k is an indicator variable that equals +1 if the k^{th} trade is a buy and -1 if the k^{th} trade is a sell, M_k is the quote midpoint immediately before the k^{th} trade, M_{IS} is the quote midpoint at the time of fund manager's request (interpreted as before the start of the trading day), P_{VWAP} is the weighted average trade price over the trading day as weighted by the volume of each trade, and P_C is the last trade price of the trading day. The total cost of trading over all trades $k = 1, 2, \ldots, K$ that take place before the deadline is

$$TCT_{T+1}^{m} = \sum_{k=1}^{K} C_k^m + I_{IS}\left(M_C - M_{IS}\right)\left(N - n_{T+1}\right), \tag{22}$$

where I_{IS} is an indicator variable that equals +1 under Implementation Shortfall and 0 otherwise and M_C is quote midpoint at the close of the trading day. The second term is an extra opportunity cost component that is part of Implementation Shortfall (but not the other metrics), which measures the foregone profits (or losses) on the quantity underfill (or overfill) $\left(N - n_{T+1}\right)$.

I incorporate the possibility that the fund manager may be informed. Let s be the fund manager's private signal prior to the trading day about the terminal value of the security and let $s - v_1$ be the difference between that signal and the time 1 public value of the security. Let $v_{T+1} - v_1$ be the cumulative innovation in the pubic value of the security from time 1 to time T+1. Define the fund manager's information ρ as the correlation between signal difference $s - v_1$ and the cumulative innovation $v_{T+1} - v_1$. For example, if the fund manager's information $\rho = 0$, then the fund manager is uninformed. Conversely, if the fund manager's information $\rho = 1$, then the fund manager is perfectly informed.

When the fund manager is unformed, the public value innovations are equally likely to be positive or negative. Conversely, if the fund manager is truly informed and chooses to request buying, then it must have been because the fund manager's signal reflected good news (i.e., the signal difference is positive $s - v_1 > 0$). Under perfect information $(\rho = 1)$ the fund manager's good news signal must have been correct, and so in this case, the public value innovations are strictly positive.[17]

Like the simple binomial model, the rich simulation is a model of a pure, open limit order market. Unlike the previous model, the initial limit order book is *not* necessarily empty.

The buy-side trader has a rich set of orders to choose from. Each period either a market order or a limit order may be submitted. A limit order may have *any* limit price as selected from a discrete price grid with penny increments. A market or limit order may be for *any* integer amount of round lots. At any time you may simultaneously cancel an unexecuted limit order and submit an updated limit or market order.

2.2 Model Calibration

In the model, I have calibrated the actions of other traders to real-world summary statistics based on order data. In 2000 the Securities and Exchange Commission (SEC) adopted Rule 605[18] that mandated the disclosure of monthly summary statistics by each exchange (or other market

---

[17] A fund manager who is truly informed and gets a signal of bad news which leads to a sell request is incorporated implicitly as the mirror image of a signal of good news leading to a buy request. In the mirror image case under perfect information $(\rho = 1)$ the fund manager's signal must have been correct and so public value innovations would be strictly negative.

[18] It was formerly named Rule 11Ac1-5.

center) for each individual stock based on order data. I select the Better Alternative Trading System (BATS) BZX exchange to calibrate to because its market structure is a pure, open limit order book just like my model. As a starting point, I calibrated other traders orders using the Rule 605 data for Microsoft for the month of December 2011.

Table I shows the calibrated inputs to the simulation. Panel A shows the probability of various order types. The Rule 605 data shows that market orders and marketable limit orders represent 61.5% of all Microsoft orders on BATS. So I assigned 30.75% to market buys and 30.75% market sells. Conversely, non-marketable limit orders represent 38.5% of all Microsoft orders on BATS. So I assigned 19.25% for non-marketable limit buys and 19.25% for non-marketable limit sells.

Panel B shows the probability of order size by order type. The Rule 605 data reports the total shares for market orders and marketable limit orders. I divided the monthly figure by 21 trading days in December 2011, by 20 intervals per day, and by 100 shares per round lot to obtain 230.4 round lots per 20-minute period. I assigned probabilities to various market order sizes to obtain an average size of 230.4 round lots. Similarly, the total shares for non-marketable limit orders were divided by 21 trading day, by 20 intervals per day, and by 100 shares per round lot to obtain 276.7 round lots per 20-minute period. I assigned probabilities to various (non-marketable) limit order sizes to obtain an average size of 276.7 round lots. Implicitly, the average calibrated order in the simulation represents the *aggregate* order quantity over a 20-minute interval.

Panel C shows the probability of price changes. I selected intraday Trade and Quote (TAQ) data for Microsoft in December 2011 and sampled trade prices every 20 minutes. To determine the intraday volatility, I computed the mean absolute deviation of the 20-minute prices as 0.033. I assigned probabilities to public value innovations such that the mean absolute deviation was 0.033.

Finally, out of all Microsoft non-marketable limit orders on BATS, 0.0% were inside-the-quote, 87.4% were at-the-quote, and 12.6% were behind-the-quote. To approximately capture this distribution, I assigned a probability of non-negative limit price deviation to be 0.0%, the probability of -$0.01 limit price deviation to be 87.4%, and the probability of a -$0.02 and -$0.03 limit price deviation to add up to 12.6%.

2.3 The First Stage of Analysis

The rich simulation model provides a wide variety of state variables that may be relevant to trading algorithms. In the first stage of analysis, I brainstormed a wide variety of trading algorithms. Specifically, I tested 46 trading algorithms which exhibited the following types of variations:

- Submitted limit buys with varying degrees of price aggressiveness relative to the bid (e.g., a penny behind the bid, equal to the bid, a penny above the bid, etc.) and with varying cancellation and resubmit policies (e.g. no cancellation, cancel and resubmit every 4 periods, cancel and resubmit every 8 periods, etc.),
- Submitted market buys with varying sizes (e.g., a few large orders, more smaller orders, many tiny orders every period, etc.),
- Opportunistically submitted additional market buys when the spread was smaller than usual (e.g., when the spread is less than 8 cents, when the spread is less than 6 cents, when the spread is less than 4 cents, etc.),
- Submitted limit orders with dynamic aggressiveness relative to the bid or relative to the midpoint and ending with a market buy
- Submitted limit orders with varying time-of-day algorithms (more front-loaded, even, or more back-loaded) and ending with a market buy.

The algorithms above were tested on a wide range problems. I varied the requested buy amount as a percentage of daily volume (10%, 50%, 100%), the deadline (2 periods, 10 periods, and 20 periods), the failure penalty (U = 0.00 and U = 0.50), fund manager risk aversion (none = 0, medium = 100, high = 1,000), fund manager information (none = 0, medium = 0.50, and high = 1.00), and performance metric (effective spread, implementation shortfall, volume-weighted average price, and closing price). Combining all of the variations, I tested 972 problems.

For each iteration of the simulation, there are four random variables that must be updated for each of the 20 time periods, plus a few more one-time random variables, for a total of 87 random variables. For each problem, I tested each algorithm with 500 iterations of the simulation and computed the average disutility over all iterations. I repeated this process for all 46 algorithms and identified the winning algorithm with the lowest average disutility. I determined the winning algorithm for all 972 problems.

2.4 The Second Stage of Analysis

In the second stage of analysis, I combined all of the winning algorithms from the first stage into a single combined algorithm. The combined algorithm is very general and flexible. It is controlled by 19 parameters that can reproduce all of the winning algorithms from the first stage of analysis as special cases. The combined algorithm encompasses a vastly larger space of possible algorithms than the original first-stage algorithms and the 19 parameters can be precisely fine-tuned to achieve the optimal result.

Specifically, the combined algorithm allows: (1) limit order only algorithms, (2) market order only algorithms, and (3) algorithms that start with limit orders and then switch to market orders when a certain number of periods remain. For the third algorithm, the switch point (in periods remaining) is an adjustable parameter.

Let x_t be the number of round lots submitted in period t. The combined algorithm allows two functional forms for number of round lots to submit

$$x_t = \begin{cases} Round\left\{\beta_{T-t} N\right\} & \text{when } I_{function} = 1 \\ Round\left\{\gamma_{T-t}\left(N - n_{t-1}\right)\right\} & \text{when } I_{function} \neq 1, \end{cases} \tag{23}$$

where β_{T-t} is the submission rate when $T-t$ periods remain under the first functional form, γ_{T-t} is the submission rate when $T-t$ periods remain under the second functional form, n_{t-1} is the cumulative number of round lots purchased by the prior period $t-1$, $I_{function}$ is a dummy parameter that selects between the two functional forms, and the continuous quantity x_t is rounded to the nearest integer.

The sequence of first function submission rates $\beta_T, \beta_{T-1}, \ldots, \beta_2, \beta_1$ is specified by three parameters as a quadratic function of the periods remaining $T-t$. Similarly, the sequence of second function submission rates $\gamma_T, \gamma_{T-1}, \ldots, \gamma_2, \gamma_1$ is specified by another three parameters as a quadratic function of the periods remaining $T-t$. Thus, these two sequences describe smooth submission rates over time that have a great deal of flexibility in functional form based on intercept, slope, and curvature parameters.

The first stage of analysis found that algorithms that use limit orders were sometimes based on the first functional form and sometimes on the second, but market orders only algorithms were always based on the second functional form. Intuitively, market orders execute in full each period, so only the quantity remaining to be obtained $N - n_{t-1}$ is relevant the next period's submission rate. By contrast, limit orders can have multiple orders outstanding that could execute any moment, so the target quantity N may determine the relevant quantity scale. Thus, I restrict market order only algorithms to the second functional form.

Two additional parameters allow an extra quantity in the first period and/or an extra quantity in the last period. The extra quantity in the beginning is useful as a way to execute the initial quantity available on the limit order book or alternatively to start with a large nonmarketable limit order algorithm, which then switches to gradually cancelling any unexecuted limit orders and resubmitting them at updated prices. An extra quantity at the end is useful to finish off limit order algorithms by placing a market order on the last period for whatever the remaining quantity is.

The sequence of limit order prices is based on the current bid price plus a smooth function of price aggressiveness relative to the current bid price. Three parameters allow the price aggressiveness function to be a quadratic function of the periods remaining $T-t$. This function on the real number line is rounded to the nearest penny price on a decimal price grid and then added to the current bid price. For example, if the price aggressiveness function rounded to the near penny is +1 cent, then the limit price will be the current bid price + 1 cent. Another parameter controls the frequency of cancelling old stale limit orders and resubmitting new limit orders at an updated price.

Finally, four dynamic aggressiveness parameters are added to the price function outlined above. If the algorithm is *behind* by a certain percentage amount, then the price is raised by a certain number of pennies. If the algorithm is *ahead* by a certain percentage amount, then the price is lowered by a certain number of pennies. For example, when there are 10 periods remaining out of an original 20 periods, a uniform algorithm should have obtained 50% of the requested quantity already. Suppose that only 30% has been obtained so far, which is 20% behind. If this is further behind than one parameter allows, then the price would be raised by the cents specified by another parameter.

For each problem as specified by the exogenous parameters, I solve for the *optimal* combined algorithm. I did this by using genetic algorithms to determine the values of the 19 parameters that minimize the disutility function. Genetic algorithms are an approach to optimization that is based on a genetic evolution paradigm. It is particularly useful when many of the parameters take on discrete values and thus cannot be optimized by traditional hill-climbing techniques.

In my application, genetic algorithms starts with random draws from the permitted range of the 19 parameters to form one “genetic unit.” Many genetic units are generated. Each genetic unit is evaluated for its “fitness.” This means that each trial trading algorithm is tested in the rich simulation against a limit order book and this process is repeated over and over again to determine the average disutility value based on 500 iterations of the limit order book. After all genetic units are tested, the best genetic units are kept and the worse units are discarded. The best genetic units are recombined with each other (“they have sex”) and a certain amount of random parameter mutations are added. Then the entire cycle repeats. The fitness of the new genetic units is determined and so on. The evolutionary process continues until the optimal genetic unit with the highest fitness is found. A typical computer run would take about five hours to determine the optimal trading algorithm such that continued computation yielded no more reduction in the average disutility function.

2.4 Results

For the figures show below, the default parameter values are a buy request that is 10% of average daily volume, average daily volume = 2,222 round lots (the Microsoft average for December 2011), the fund manager’s deadline $T = 20$ (corresponding to 4:00 p.m.), the failure penalty $U = 0.00$, risk aversion $A = 0$, fund manager’s information = 0, and the performance metric is implementation shortfall. Thus, the default buy request is 10%*2,222 round lots and rounded to the nearest integer to get 222 round lots.

Figures 3 - 7 show the cumulative round lots purchased over time by the optimal trading algorithm for different trading problems. Solid-filled markers indicate that a market order or marketable limit order was submitted this period. Unfilled markers indicate that a nonmarketable limit order was submitted this period.

Figure 4 shows the average cumulative round lots purchased over time by the optimal trading algorithm under different failure penalties. The solid curve is for the high failure penalty (U = 0.50) and the dashed curve is for no failure penalty (U = 0.00). Under a high failure penalty, the optimal trading algorithm is to submit nonmarketable limit orders for first nine periods and then switch to submitting marketable orders for ten more periods. The solid curve shows the cumulative lots purchased over time *on average* over many scenarios (iterations). But irrespective of whether the initial set of limit orders buys few lots or many lots in a given scenario, the subsequent marketable orders speed up or slow down, so that this algorithm purchases *exactly* the requested 222 round lots by the end. By contrast, under a zero failure penalty the optimal trading algorithm is to submit strictly nonmarketable limit orders for all 20 periods. Since there is no penalty for failing to obtain the requested quantity, every single limit order must earn a negative cost (a.k.a., a profit). In general, a larger quantity of limit order lots are submitted than execute, Thus, great care is taken to avoid going over the requested 222 round lots, since there is a large penalty for going over the requested amount. As a consequence, the average outcome is purchasing 177 round lots by the end.

Figure 5 shows the average cumulative round lots purchased over time by the optimal trading algorithm under different portfolio manager's degree of information ranging from 0.00

(uninformed) to 0.20 (relatively high info). The dashed curve shows that when the portfolio manager is uninformed (info = 0.00), the optimal trading algorithm is to submit limit orders so as to purchase 97 lots by period 19 and then a final market order at the end to get the remaining quantity. When the degree of information is 0.05, the limit orders become more front-loaded so as to purchase 146 lots by period 17 and purchase the rest with market orders. As the degree of information increases, the optimal algorithm becomes more front-loaded. This makes intuitive sense given we are only considering buying requests,[19] because a higher degree of information implies a forecast that the price is more likely to increase and front-loading buys a at a lower price before the price has risen. When the degree of information is 0.15, the limit orders become so front-loaded as to purchase 139 lots by period 6 and purchase the rest with market orders by period 15. Finally, when the degree of information is 0.20 (or anything higher), the optimal algorithm is market buy for the full amount in the first period, which is the ultimate extreme in front-loading.

Figure 6 shows the cumulative round lots purchased over time by the optimal trading algorithm under different risk aversions. The large-dashed curve shows that when the portfolio manager is risk neutral, the optimal trading algorithm is to submit limit orders so as to purchase 97 lots by period 19 and then a final market order at the end to get the remaining quantity. The small-dashed curve (solid curve) shows when the portfolio manager has medium (high) absolute risk aversion, the optimal algorithm becomes more front-loaded (even more front-loaded). This makes intuitive sense, because buying lots early in the day reduces the standard deviation of trade prices relative to the implementation shortfall benchmark, which is the quote midpoint before the start of the trading day.

[19] Selling requests are incorporated indirectly as the mirror image of buying requests.

Figure 7 shows the cumulative round lots purchased over time by the optimal trading algorithm under different performance metrics. The solid curve shows the default case that we have considered above, namely implementation shortfall, which optimally involves submitting limit orders so as to purchase 97 lots by period 19 and then a final market order at the end to get the remaining quantity. The other three performance metrics, Volume Weighted Average Price, Effective Spread, and Closing Price, lead to mildly more front-loaded algorithms that start with limit orders and then switch to marketable orders to get the full amount by the end.

Table II shows the optimal trading algorithm by failure penalty, risk aversion, and performance metric. Panel A shows when the failure penalty is zero, which represents the most purely opportunistic fund manager. Strikingly, the optimal algorithm disutility is always negative. This because the optimal algorithm in each case involves limit buys only – no market buys. If there is no failure penalty for failing to get the requested quantity, then the only trades that you want to do are those that incur a *negative* cost of trading, meaning a profit. In other words, you want every single trade to make a profit. The last column indicates whether the optimal algorithm involves interior optimum values (i.e., not corner solutions) for the dynamic aggressiveness parameters. In most cases (10 out of 12), they do.

Panel B shows when the failure penalty is high (0.50). In every case, the optimal algorithm involves limit buys early, followed by market buys later. The dynamic aggressiveness results are a little more mixed with the majority (7 out of 12) using dynamic aggressiveness.

I also compare these optimal algorithms to two benchmark algorithms from the existing literature. Recall that the market order only benchmark algorithm is to place equal-sized market orders over the entire trading horizon and the limit-then-market benchmark algorithm is to place equal-sized limit orders first followed by a market order to get any remaining amount in the last period. Table II reports the average disutility obtained by the optimal algorithm and the two benchmark algorithms. The optimal algorithms beat the market order only algorithm in all cases, beat the limit-then-market algorithm in 21 cases, and tie the limit-then-market algorithm in 3 cases. In these cases, the optimal strategy *is* equal-limits-then-market algorithm. In summary, the optimal trading algorithm frequently beats (or at least ties) the benchmark trading algorithms from the existing literature.

Table III shows the optimal trading algorithm by fund manager's information, deadline, and performance metric. Panel A shows when the fund manager's information is zero. When the deadline is very short (2 periods), then the optimal trading algorithm is very aggressive with large market buys. When the deadline is medium or long (10 or 20 periods), then the optimal trading algorithm is mainly limit buys, but with market buys at the end to get the requested amount. All of the optimal trading algorithms in Table III use dynamic aggressiveness.

Panel B shows when the fund manager's information is high (0.50). As before when the deadline is very short, it is best to go with large market buys. However, when the deadline is medium or long and the performance metric is *not* effective spread, then it is also best to submit large market buys. The intuitive explanation for this has three steps: (1) the fund manager has a private signal of good news leading to a buy request, (2) statistically the fund manager is correct leading to a price increase over the trading day on average, and (3) anticipating this, it is best to purchase the security early in the day before the price has risen very much.

An interesting exception to this logic occurs when the performance metric *is* effective spread. In this case, it is best to submit mainly limit buys, but with market buys at the end. The benchmark under the effective spread metric is the contemporaneous midpoint of each trade. This has the unique property that the benchmark moves upward if trades are later in the day at higher prices. So under effective spread, there is no measured cost penalty for trading later in the day and it is best to spread out the trades over time. By contrast, the benchmark in the other three performance metrics (the request midpoint, the volume-weighted average price, and the closing price) are unaffected by the timing of trades. So under these three metrics, it is best to aggressively front-load the trading when you anticipate that the price will rise on average.

Finally, Table III reports the average disutility obtained by the optimal algorithm compared to the two benchmark algorithms. The optimal algorithms beat the market order only algorithm in all cases, beat the equal-limit-then-market algorithm in 15 cases, and tie the equal-limit-then-market algorithm in 9 cases because the equal-limit-then-market algorithm *is* the optimal algorithm in those cases. In summary, the optimal trading algorithms frequently beat (or at least tie) the benchmark trading algorithms from the existing literature.

3. Conclusion

I determine the optimal trading strategy for an institutional trader who wants to purchase a large number of shares over a fixed time horizon. First, I consider when a multi-period trader can submit limit orders as well as market orders. I develop a simple binomial model where limit orders either execute or not and solve it analytically. I find that the optimal sequence of limit orders involves dynamic aggressiveness. That is, if a given limit order executes (or not), then the next limit order optimally has a slightly less (more) aggressive price. I find that this trading strategy frequently beats (or at least ties) the benchmark trading strategies from the existing literature. Second, I develop an elaborate model of a multi-period trading algorithm that allows multiple orders over multiple periods to be submitted to a dynamic limit order book exchange. I calibrate the simulation to real-world summary statistics based on order data. For each trading problem, I use genetic algorithm optimization to numerically determine the optimal trading algorithm over a large number of simulated trials. I find that the optimal trading algorithm typically involves dynamic aggressiveness. Further, I find that if the fund manager is opportunistic, then the optimal algorithm involves only limit orders with low price aggressiveness. Conversely if the fund manager is committed, then limit orders should be followed by market orders at the end. I find that if the fund manager is informed and not using effective spread to measure the cost of trading, then market orders should be front-loaded in time. Conversely, if effective spread is used or the fund manager is uninformed, then less aggressive orders should be spread evenly over time. I find that these trading algorithms frequently beat (or at least tie) the benchmark trading algorithms from the existing literature.

Appendix

Proof of Proposition 1. When $k < 3d$, the optimal limit price aggressiveness is obtained by substituting (1) into (2) and taking the derivative with respect to b. The second order condition is positive and thus the objective function is minimized. The probability of execution is obtained by substituting (3) in place of $b - v$ in (1). The expected disutility is obtained by substituting (3) and (4) into (2). When $k = 3d$, substitute $3d$ in place of k in (3), (4), and (5) to obtain $a = d$, $p = 1$, and $E[D] = d$. By definition the price aggressiveness can't go above d, so for any value of $k > 3d$ the solution is constrained to the corner solution $a = d$, $p = 1$, and $E[D] = d$. *Q.E.D.*

Proof of Proposition 2. The optimal limit price aggressiveness is obtained by substituting (6), (10), and (11) into (9) and taking the derivative with respect to b. The second order condition is positive and thus the objective function is minimized. The probability of execution is obtained by substituting (12) in place of $b_{t,n} - v_t$ in (6). Equation (14) for the expected total cost of trading on future trades at the (t,n) node is the probability weighted average of the expected total cost of trading on future trades at $(t+1,n+1)$ node and the $(t+1,n)$ node, where the incremental cost of trading $b_{t,n}-v_t$ is added when the current order executes. Equation (15) for the expected number of units to be purchased in future trades at the (t,n) node is the probability weighted average of the expected number of units to be purchased in future trades at $(t+1,n+1)$ node and the $(t+1,n)$ node, where the number of units is incremented by one unit when the current order executes. *Q.E.D.*

Combined Proof of Propositions 3 and 4. The first step of the combined proof is to show that the probability of execution at any node where the trader is not done $(n<N)$ is a weighted average of the following two connected nodes. The proof is based on the connection between the *f* and *h* variables at a given node and the *f* and *h* variables at the following connected notes. Substituting (12) into (13) and then twice substituting (14) and (15) into the resulting equation yields

$$p_{t,n}=\left(\frac{1}{4d}\right)\left(\begin{array}{l}\left[p_{t+1,n}\left(b_{t+1,n}-v_{t+1}+f_{t+2,n+1}\right)+\left(1-p_{t+1,n}\right)f_{t+2,n}\right]\\-\left[p_{t+1,n+1}\left(b_{t+1,n+1}-v_{t+1}+f_{t+2,n+2}\right)+\left(1-p_{t+1,n+1}\right)f_{t+2,n+1}\right]\\+k\left(\begin{array}{l}1+\left[p_{t+1,n+1}\left(1+h_{t+2,n+2}\right)+\left(1-p_{t+1,n+1}\right)h_{t+2,n+1}\right]\\-\left[p_{t+1,n}\left(1+h_{t+2,n+1}\right)+\left(1-p_{t+1,n}\right)h_{t+2,n}\right]\end{array}\right)+d\end{array}\right). \tag{17}$$

Rearranging (17) yields

$$p_{t,n}=\left(\frac{1}{4d}\right)\left(\begin{array}{l}p_{t+1,n}\left(b_{t+1,n}-v_{t+1}-\left[f_{t+2,n}-f_{t+2,n+1}+k\left(\left(1+h_{t+2,n+1}-h_{t+2,n}\right)\right)\right]\right)\\-p_{t+1,n+1}\left(b_{t+1,n+1}-v_{t+1}-\left[f_{t+2,n+1}-f_{t+2,n+2}+k\left(1+h_{t+2,n+2}-h_{t+2,n+1}\right)\right]\right)\\+\left[f_{t+2,n}-f_{t+2,n+1}+k\left(\left(1+h_{t+2,n+1}-h_{t+2,n}\right)\right)\right]+d\end{array}\right). \tag{18}$$

Substituting from (12) into (18) yields

$$p_{t,n} = \left(\frac{1}{4d}\right)\begin{pmatrix} p_{t+1,n}\left(a_{t+1,n} - \left[2a_{t+1,n} + d\right]\right) \\ -p_{t+1,n+1}\left(a_{t+1,n+1} - \left[2a_{t+1,n+1} + d\right]\right) \\ +\left[2a_{t+1,n} + d\right] + d \end{pmatrix}. \tag{19}$$

Substituting from (13) into (19) yields

$$p_{t,n} = \left(\frac{1}{4d}\right)\begin{pmatrix} p_{t+1,n}\left(d\left(2p_{t+1,n} - 1\right) - d\left(4p_{t+1,n} - 1\right)\right) \\ -p_{t+1,n+1}\left(d\left(2p_{t+1,n+1} - 1\right) - d\left(4p_{t+1,n+1} - 1\right)\right) \\ +d\left(4p_{t+1,n} - 1\right) + d \end{pmatrix}. \tag{20}$$

Rearranging (20) yields equation (16), in which the probability of execution at any node where the trader is not done $\left(n < N\right)$ is a weighted average of the following two connected nodes. Equation (16) will yield dynamic aggressiveness in the Limit Order Zone once it is proven that up-nodes have higher probabilities and down-nodes have lower probabilities.

The next step is to define the Market Order / Maximum Aggressiveness Zone for both the *a* and *p* binomial trees at the set of all nodes where there is non-positive slack time (i.e., the same or fewer time periods left than remaining units to be purchased $T + 1 - t \leq N - n$). At all terminal nodes, there are no future trades, so $f_{T+1,n} = 0,$ and $h_{T+1,n} = 0.$ Substitute these equations into (12) for all of the nodes on the last trading date $\left(t = T\right)$ in the Market Order / Maximum Aggressiveness Zone to get

$$a_{T,n} = \begin{Bmatrix} \frac{1}{2}\left(k - d\right) & \text{when } k < 3d \\ d & \text{when } k > 3d \end{Bmatrix}, \tag{21}$$

and substitute this into (13) to get

$$p_{T,n} = \begin{Bmatrix} \dfrac{k + d}{4d} & \text{when } k < 3d \\ 1 & \text{when } k > 3d \end{Bmatrix}. \tag{22}$$

(21) and (22) are identical to the single-period solution. All of the nodes in this zone on earlier trading dates $(t < T)$ have two following connected nodes with a probability of execution given by (22). Since all nodes are a weighted average of the following two connected nodes by equation (16), their probability of execution must also be given by (22). Inverting (13) yields

$$a_{t,n} = d\left(2p_{t,n} - 1\right). \tag{23}$$

Substituting (22) into (23) yields (21).

The next step is to define the Limit Order Zone for both the a and p binomial trees at the set of all nodes where the trader is not done $(n < N)$ and there is positive slack time (i.e., more time periods left than remaining units to be purchased $T + 1 - t > N - n$). Consider all of the nodes in the Limit Order Zone that border the Market Order / Maximum Aggressiveness Zone (i.e., have one following connection in the Market Order / Maximum Aggressiveness Zone and one following connection *not* in the Market Order / Maximum Aggressiveness Zone). Since all nodes are a weighted average of the following two connected nodes by equation (16), these nodes must have a probability of execution that is strictly less than (22) and substituting into (23) yields an optimal limit order aggressiveness strictly less than (21).

The next step is to define the Done Zone for both the a and p binomial trees at the set of all nodes where the trader is done $(n \geq N)$. For all of the nodes in the Done Zone, no more orders will be submitted and so the probability that an order will execute is zero $p_{t,n} = 0$. Consider all of the nodes in the Limit Order Zone that border the Done Zone (i.e., have one following connection in the Done Zone and one following connection *not* in the Done Zone). Since all nodes are a weighted average of the following two connected nodes by equation (16), these nodes must have a probability of execution that is strictly greater than zero and substituting into (23) must have an optimal limit order aggressiveness strictly greater than $-d$.

The final step is to consider all nodes in the Limit Order Zone. Since all nodes are a weighted average of the following two connected nodes by equation (16), then one of the following connected nodes will have a higher probability of execution and the other will have a lower probability of execution. Since the nodes in the Limit Order Zone that border the Market Order / Maximum Aggressiveness Zone must connect to the *higher* probability of execution of the Market Order / Maximum Aggressiveness Zone and since the nodes in the Limit Order Zone that border the Done Zone must connect to the *lower (zero)* probability of execution of the Done Zone, then up-nodes must have a higher probability of execution and down-nodes must have a lower probability of execution. This proof that up-nodes have higher probabilities and down-nodes have lower probabilities when combined with (16) yield that the probabilities exhibit dynamic aggressiveness in the Limit Order Zone, and by substituting into (23), their optimal limit order aggressiveness must also exhibit dynamic aggressiveness in the Limit Order Zone. *Q.E.D.*

References

Almgren, R. and Chriss, N., 1999, Value under liquidation, *Journal of Risk* 1, 61-63.

Almgren, R. and Chriss, N., 2001, Optimal execution of portfolio transactions, *Journal of Risk* 3, 5-39.

Angel, J., 1995, Limit Versus Market Orders, Unpublished paper, Georgetown University.

Bertsimas, D. and A. Lo, 1998, Optimal control of execution costs, *Journal of Financial Markets* 1, 1-50.

Bloomfield, R., M. O'Hara, and G. Saar, 2005, The 'Make or Take' Decision in an Electronic Market: Evidence on the Evolution of Liquidity, *Journal of Financial Economics* 75, 165-199.

Buti, S. and B. Rindi, 2013, Undisclosed Orders and Optimal Submission Strategies in a Limit Order Market, *Journal of Financial Economics* 109, 797-812.

Colliard, J. and Foucault, T., 2012, Trading Fees and Efficiency in Limit Order Markets, *Review of Financial Studies* 25, 3389-3421.

Conrad, J., Johnson, K., and S. Wahal, 2003, Institutional Trading and Alternative Trading Systems, *Journal of Financial Economics* 70, 99-134.

Domowitz, I., Glen, J. and A. Madhavan, 2001, Liquidity, Volatility, and Equity Trading Costs Across Countries and Over Time, *International Finance* 4, 221-255.

Foucault, T., 1999, Order Flow Composition and Trading Costs in a Dynamic Limit Order Market, *Journal of Financial Markets* 2, 99-134.

Foucault, T., Kadan, O., and E. Kandel, 2005. Limit Order Book as a Market for Liquidity, *Review of Financial Studies* 18, 1171-1217.

Goettler, R., C. Parlour, and U. Rajan, 2005, Equilibrium in a Dynamic Limit Order Book, *Journal of Finance*, 60, 2149-2192.

Goettler, R., Christine P., and U. Rajan, 2009. Informed traders and limit order markets, *Journal of Financial Economics* 93, 67-87.

Handa, P., and R. Schwartz, 1996, Limit Order Trading, *Journal of Finance* 51, 1835-61.

Harris, L., 1998, Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems, *Financial Markets, Institutions, and Instruments* 7, 1-75.

Harris, L., and J. Hasbrouck, 1996, Market versus Limit Orders: the SuperDOT Evidence on Order Submission Strategy, *Journal of Financial and Quantitative Analysis* 31, 212-31.

Huberman, G. and W. Stanzl, 2005, Optimal liquidity trading, *Review of Finance* 9, 165–200.

Irvine, P, Lipson, M. and A. Puckett, 2007, Tipping, *Review of Financial Studies* 20, 741-768.

Keim, D. and A. Madhavan, 1995, Anatomy of the Trading Process: Empirical Evidence on the behavior of institutional traders, *Journal of Financial Economics* 37, 371-398.

Keim, D. and A. Madhavan, 1996, The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects, *Review of Financial Studies* 9, 1-36.

Keim, D. and A. Madhavan, 1997, Transaction Costs and Investment Style: An Inter-exchange Analysis of Institutional Equity Trades, *Journal of Financial Economics* 46, 265-292.

Lakonishok, J., Schleifer, A., and R. Vishny, 1992, The impact of institutional trading on stock prices, *Journal of Financial Economics* 32, 23-43.

Obizhaeva, A. and Wang, J., 2013, Optimal trading strategy and supply/demand dynamics, *Journal of Financial Markets* 16, 1-32.

Parlour, C., 1998, Price Dynamics in Limit Order Markets, *Review of Financial Studies* 11, 789-816.

Parlour, C. and D. Seppi, 2003, Liquidity-Based Competition for Order Flow, *Review of Financial Studies* 16, 301-343.

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

Vayanos, D., 2001, Strategic trading in a dynamic noisy market, *Journal of Finance* 56, 131-171.

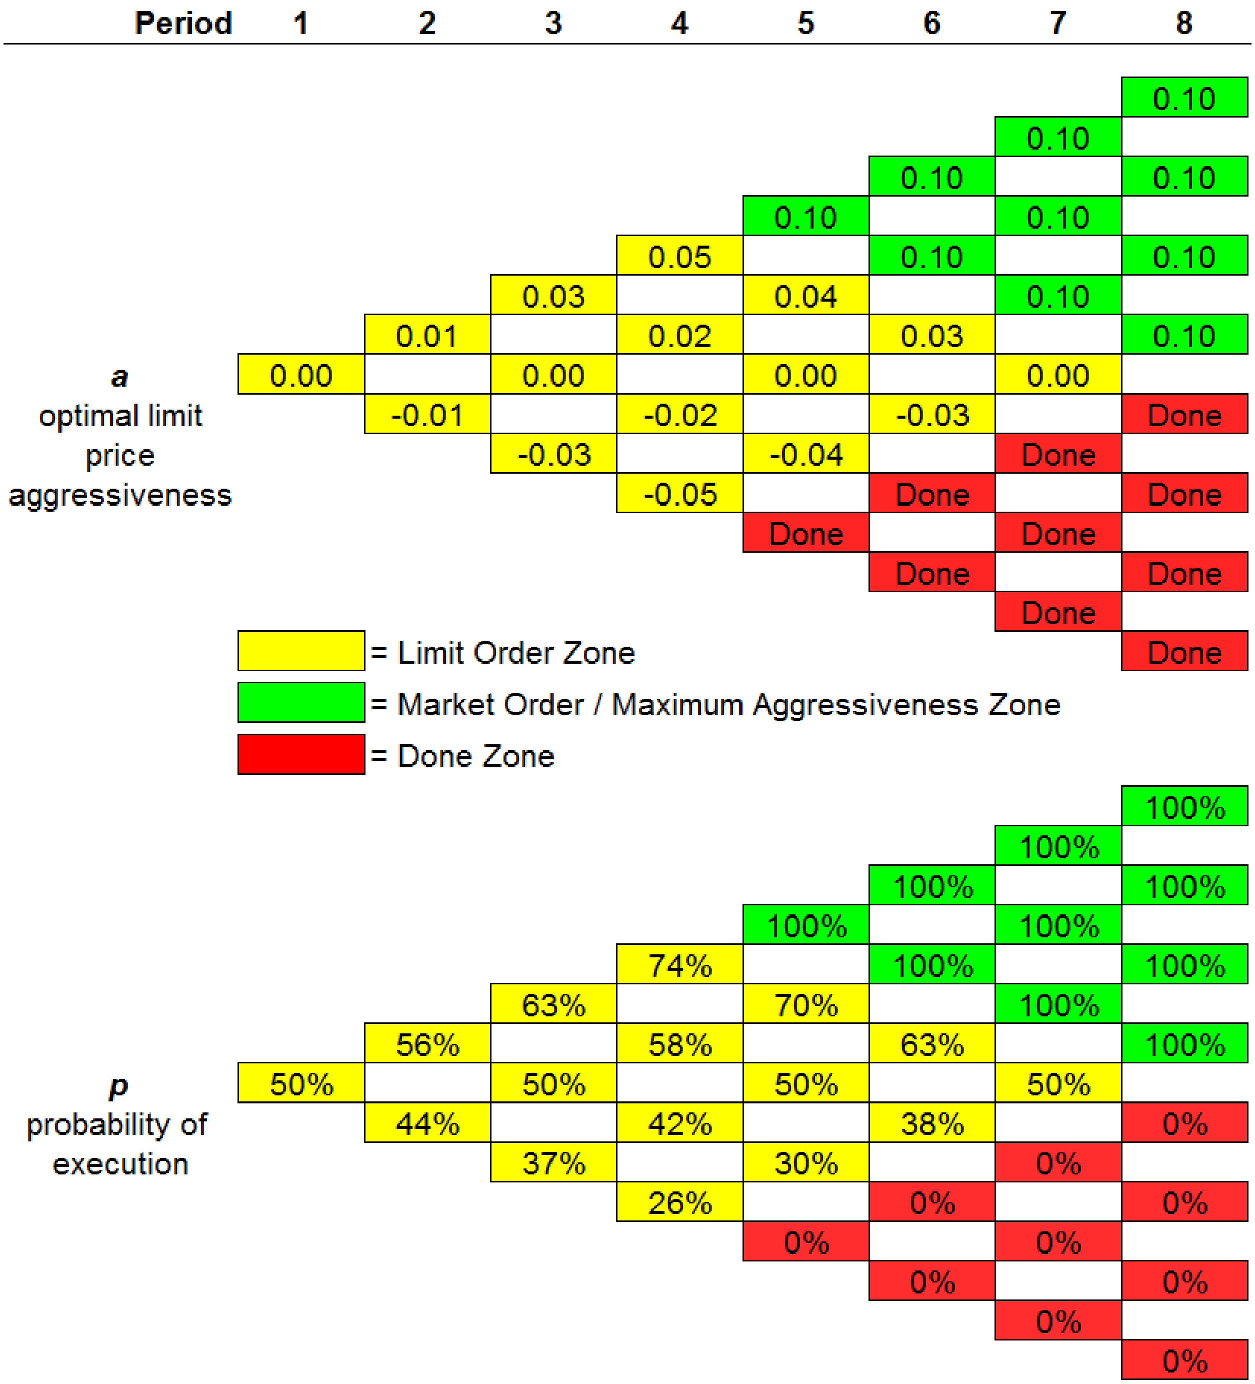

Figure 1. Binomial Trees for *a* and *p* when the Failure Penalty is High.

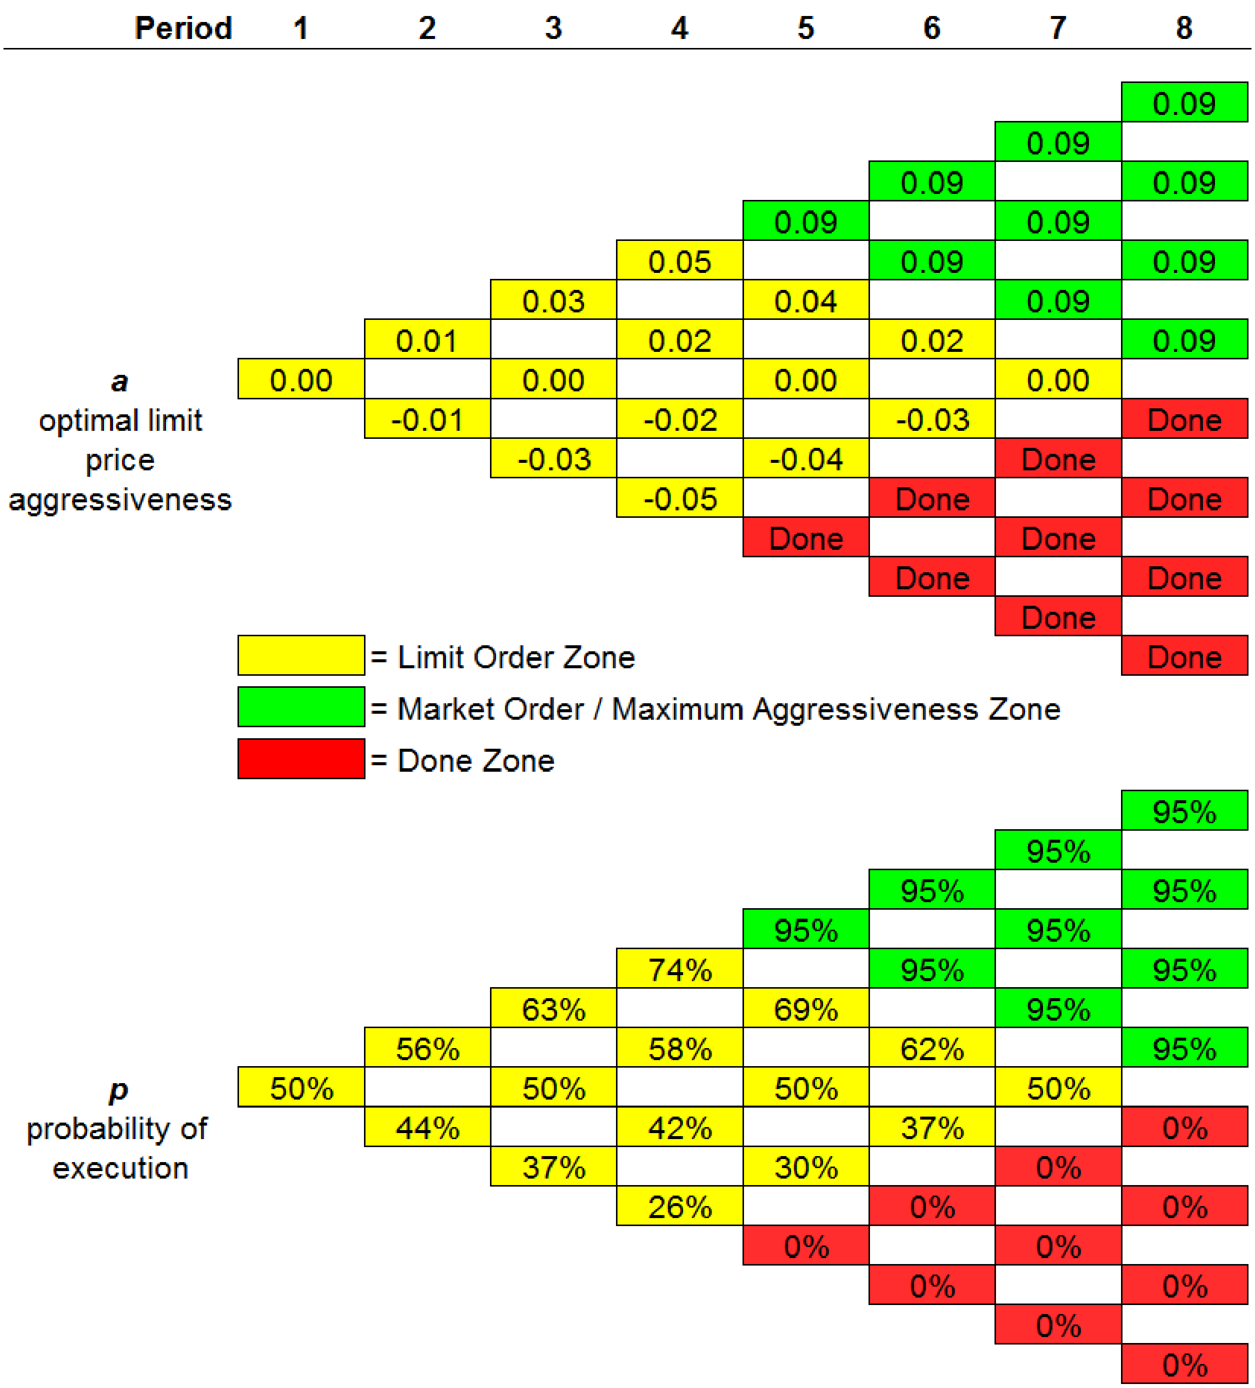

Figure 2. Binomial Trees for *a* and *p* when the Failure Penalty is Low.

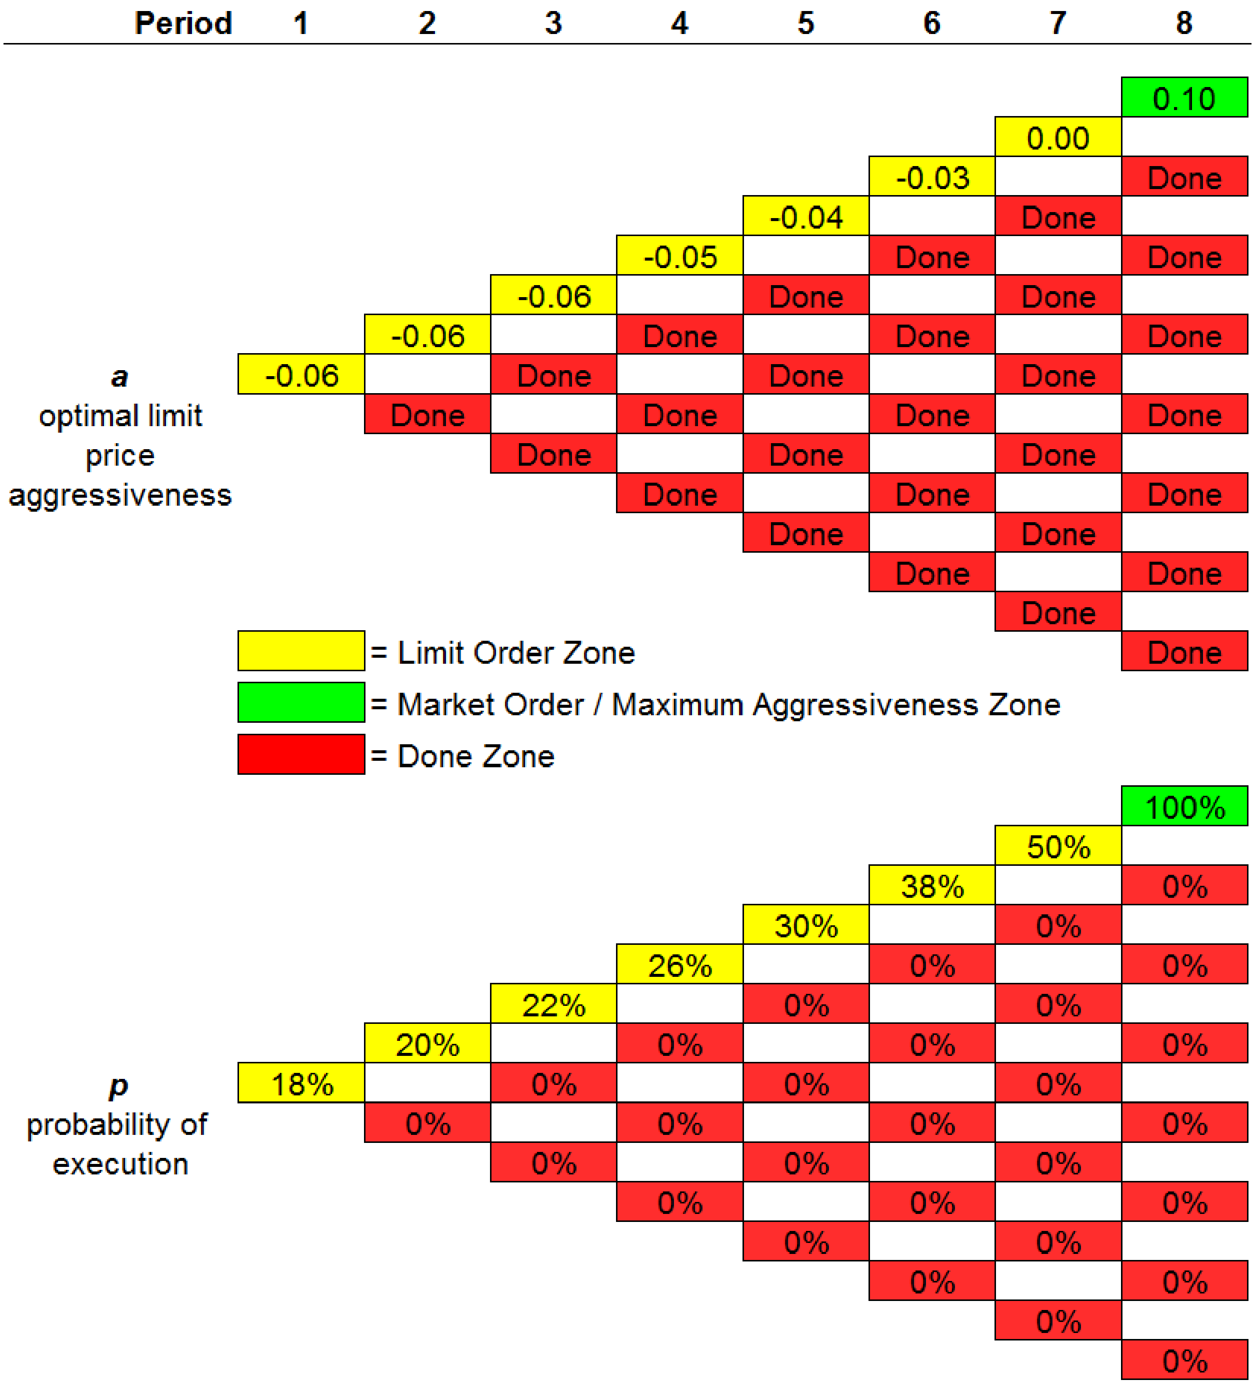

Figure 3. Binomial Trees for *a* and *p* when $N = 1$ and the Failure Penalty is High.

Table I
Calibration of Simulation Input Parameters

The simulation input parameters are calibrated to SEC Rule 605 mandated disclosure of summary statistics by exchange by stock. Summary statistics for Dec 2011 are divided by 21 trading days and by 20 intervals to scale to 20 minute intervals. NYSE TAQ data is sampled every 20 minutes to determine price volatility at 20 minute intervals.

Panel A: Probability of order types

Order Type	Probability
Market buy	30.75%
Market sell	30.75%
Limit buy	19.25%
Limit sell	19.25%
No order	0.00%
Total	100.00%

Panel B: Probability of Order Size by Order Type

Order Size	Probability of Market Order	Probability of Limit Order
10	3.0%	0.5%
20	3.0%	1.0%
30	3.0%	1.0%
40	4.0%	1.0%
50	4.0%	1.5%
100	19.0%	20.0%
200	20.0%	20.0%
300	20.0%	20.0%
400	14.0%	20.0%
500	10.0%	15.0%
Total	100.0%	100.0%
Average Size	230.4	276.7

Panel C: Probability of Price Changes

Price Changes	Probability of Public Value Innovations	Probability of Limit Price Deviations
-$0.05	13.0%	0.0%
-$0.04	12.0%	0.0%
-$0.03	10.0%	4.6%
-$0.02	8.0%	8.0%
-$0.01	6.0%	87.4%
$0.00	2.0%	0.0%
$0.01	6.0%	0.0%
$0.02	8.0%	0.0%
$0.03	10.0%	0.0%
$0.04	12.0%	0.0%
$0.05	13.0%	0.0%
Total	100.0%	100.0%
Mean Absolute Deviation	0.033	

Table II

Optimal Trading Algorithm By Failure Penalty, Risk Aversion, and Performance Metric

The optimal trading algorithm is shown by failure penalty, risk aversion, and performance metric. In all cases below, the requested buy is 10% of daily volume, the deadline is 20, and the fund manager's information is zero.

Risk Aversion	Performance Metric	Optimal Strategy Disutility	Market Order Only Benchmark Disutility	Limit, Then Market Bchmrk. Disutility	Gain Over First Benchmark	Gain Over Second Benchmark	Dynamic Aggressive -ness
Panel A Failure Penalty is Zero							
0	Effective Spread	-10.6	5.6	-1.7	16.2	8.9	Yes
0	Imple. Shortfall	-8.1	7.3	2.2	15.4	10.3	Yes
0	VWAP	-4.7	6.1	0.9	10.8	5.6	No
0	Closing Price	-2.5	0.5	-0.5	3.0	2.0	Yes
100	Effective Spread	-8.8	6.3	-1.6	15.1	7.2	Yes
100	Imple. Shortfall	-11.4	8.8	4.3	20.1	15.7	Yes
100	VWAP	-10.3	6.5	1.7	16.9	12.0	No
100	Closing Price	-6.4	1.0	2.2	7.4	8.7	Yes
1000	Effective Spread	-9.1	7.3	-1.8	16.4	7.3	Yes
1000	Imple. Shortfall	-7.1	9.1	2.4	16.2	9.5	Yes
1000	VWAP	-9.8	7.3	0.6	17.1	10.5	Yes
1000	Closing Price	-6.9	1.2	-0.3	8.1	6.6	Yes
Panel B Failure Penalty is High (0.50)							
0	Effective Spread	-0.9	22.1	-0.9	23.0	0.0	Yes
0	Imple. Shortfall	3.0	23.8	3.0	20.8	0.0	Yes
0	VWAP	1.7	22.6	1.7	21.0	0.0	Yes
0	Closing Price	-7.0	17.0	0.3	24.0	7.3	Yes
100	Effective Spread	-4.2	22.2	-0.6	26.4	3.6	No
100	Imple. Shortfall	0.9	24.7	5.3	23.9	4.5	No
100	VWAP	-4.4	22.5	2.7	26.9	7.1	No
100	Closing Price	-0.5	17.3	3.5	17.8	4.0	No
1000	Effective Spread	-7.3	22.4	-0.4	29.7	6.9	Yes
1000	Imple. Shortfall	1.9	24.3	3.8	22.4	1.8	Yes
1000	VWAP	-1.3	22.4	2.0	23.7	3.2	Yes
1000	Closing Price	-2.8	16.4	1.1	19.1	3.8	No

Table III

Optimal Trading Algorithm By Fund Manager's Information, Deadline, and Performance Metric

The optimal trading algorithm is shown by fund manager's information, deadline, and performance metric. In all cases, the requested buy is 10% of daily volume, the failure penalty is 0.50, and the fund manager's risk aversion is zero.

Deadline	Performance Metric	Optimal Strategy Disutility	Only Benchmark Disutility	Market Bchmrk. Disutility	Gain Over First Benchmark	Gain Over Second Benchmark	Dynamic Aggressive -ness
Panel A Fund Manager's Information is Zero							
2	Effective Spread	4.3	5.0	5.2	0.7	0.9	Yes
2	Imple. Shortfall	4.7	5.8	7.1	1.1	2.4	Yes
2	VWAP	3.7	4.6	5.4	0.9	1.7	Yes
2	Closing Price	6.0	6.5	7.6	0.6	1.6	Yes
10	Effective Spread	-0.4	13.6	-0.4	14.0	0.0	Yes
10	Imple. Shortfall	3.5	15.1	3.5	11.7	0.0	Yes
10	VWAP	2.0	13.8	2.0	11.8	0.0	Yes
10	Closing Price	3.3	15.1	3.3	11.8	0.0	Yes
20	Effective Spread	-0.9	22.1	-0.9	23.0	0.0	Yes
20	Imple. Shortfall	3.0	23.8	3.0	20.8	0.0	Yes
20	VWAP	1.7	22.6	1.7	21.0	0.0	Yes
20	Closing Price	-7.0	17.0	0.3	24.0	7.3	Yes
Panel B Fund Manager's Information is High (0.50)							
2	Effective Spread	4.2	5.0	5.2	0.8	1.0	Yes
2	Imple. Shortfall	5.2	6.8	8.7	1.6	3.6	Yes
2	VWAP	-4.3	-3.3	-3.1	1.0	1.2	Yes
2	Closing Price	-25.5	-24.2	-25.4	1.3	0.1	Yes
10	Effective Spread	-0.3	16.1	-0.3	16.4	0.0	Yes
10	Imple. Shortfall	5.2	23.1	5.7	18.0	0.5	Yes
10	VWAP	-4.3	10.7	-3.8	15.0	0.5	Yes
10	Closing Price	-25.5	-9.3	-23.0	16.2	2.4	Yes
20	Effective Spread	-1.7	23.6	-1.3	25.3	0.4	Yes
20	Imple. Shortfall	4.0	40.4	5.5	36.4	1.5	Yes
20	VWAP	-5.3	25.6	-4.1	30.9	1.2	Yes
20	Closing Price	-28.1	-4.6	-28.1	23.5	0.0	Yes

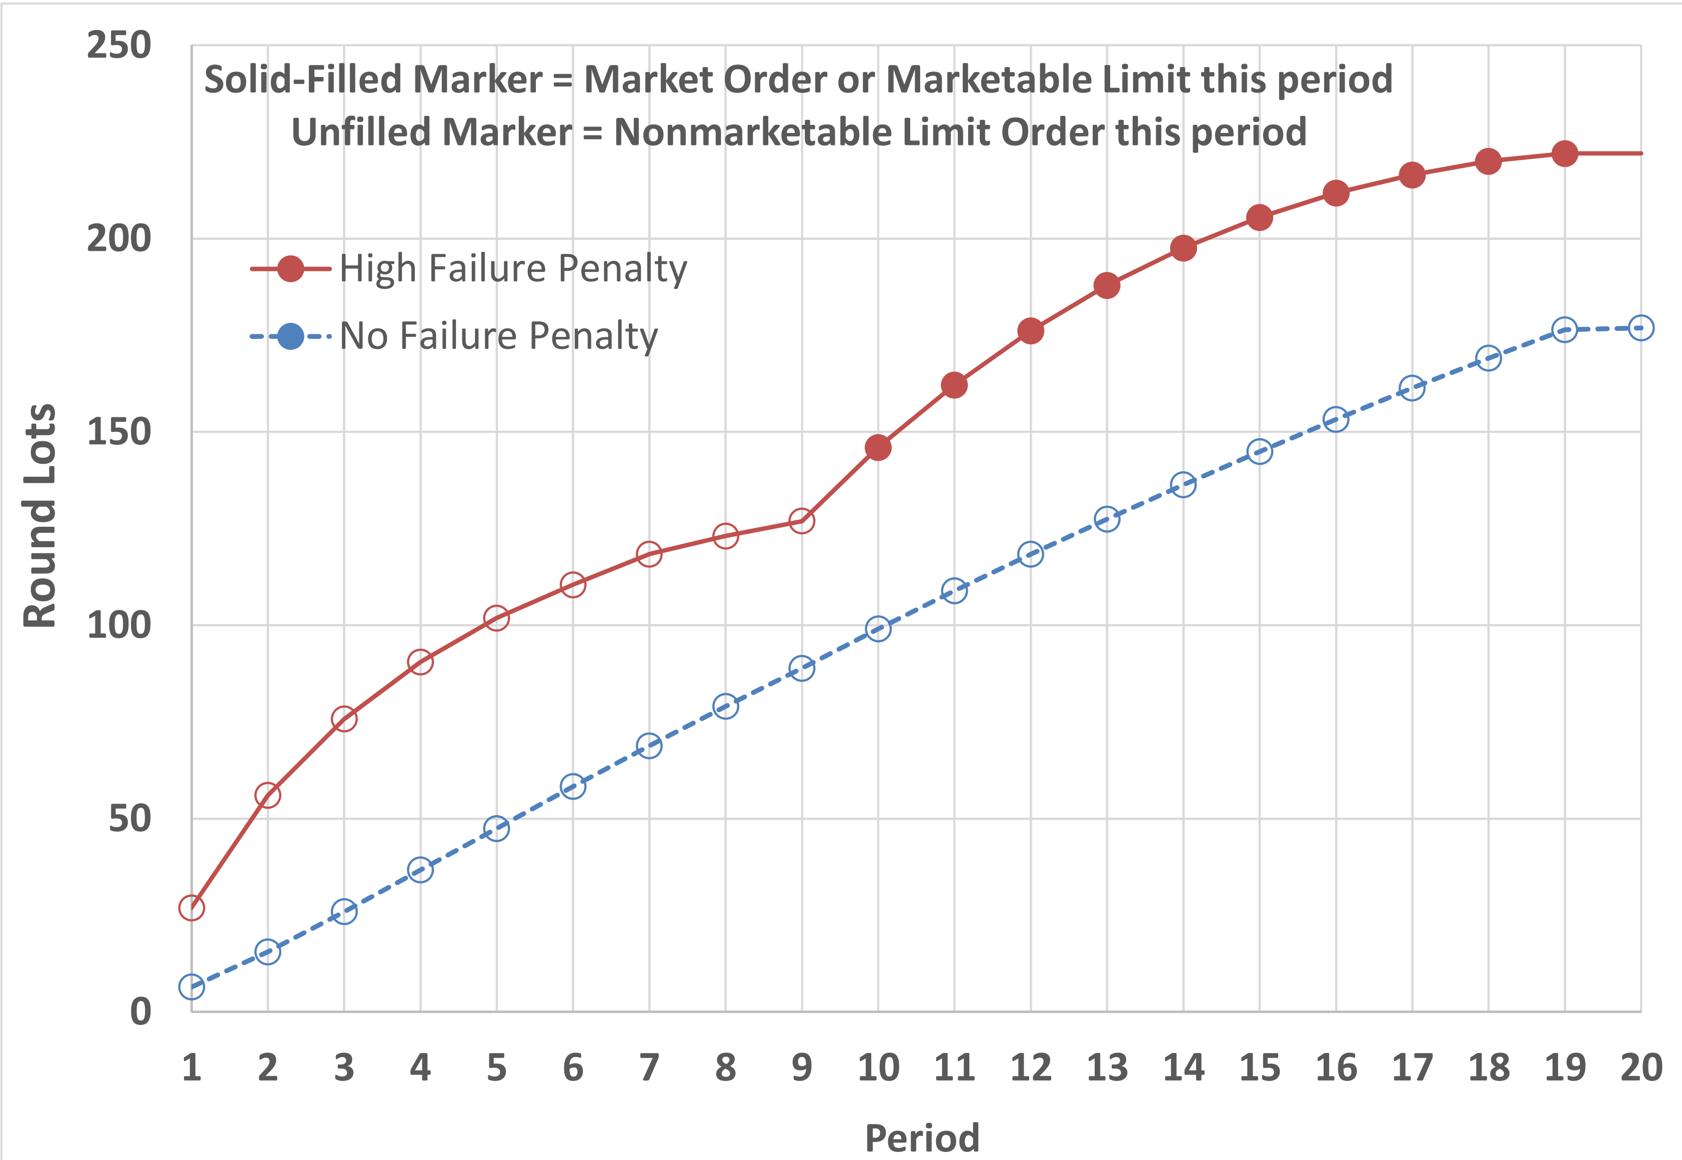

Figure 4. Average Cumulative Round Lots Bought Over Time by the Optimal Trading Strategy for Different Failure Penalties.

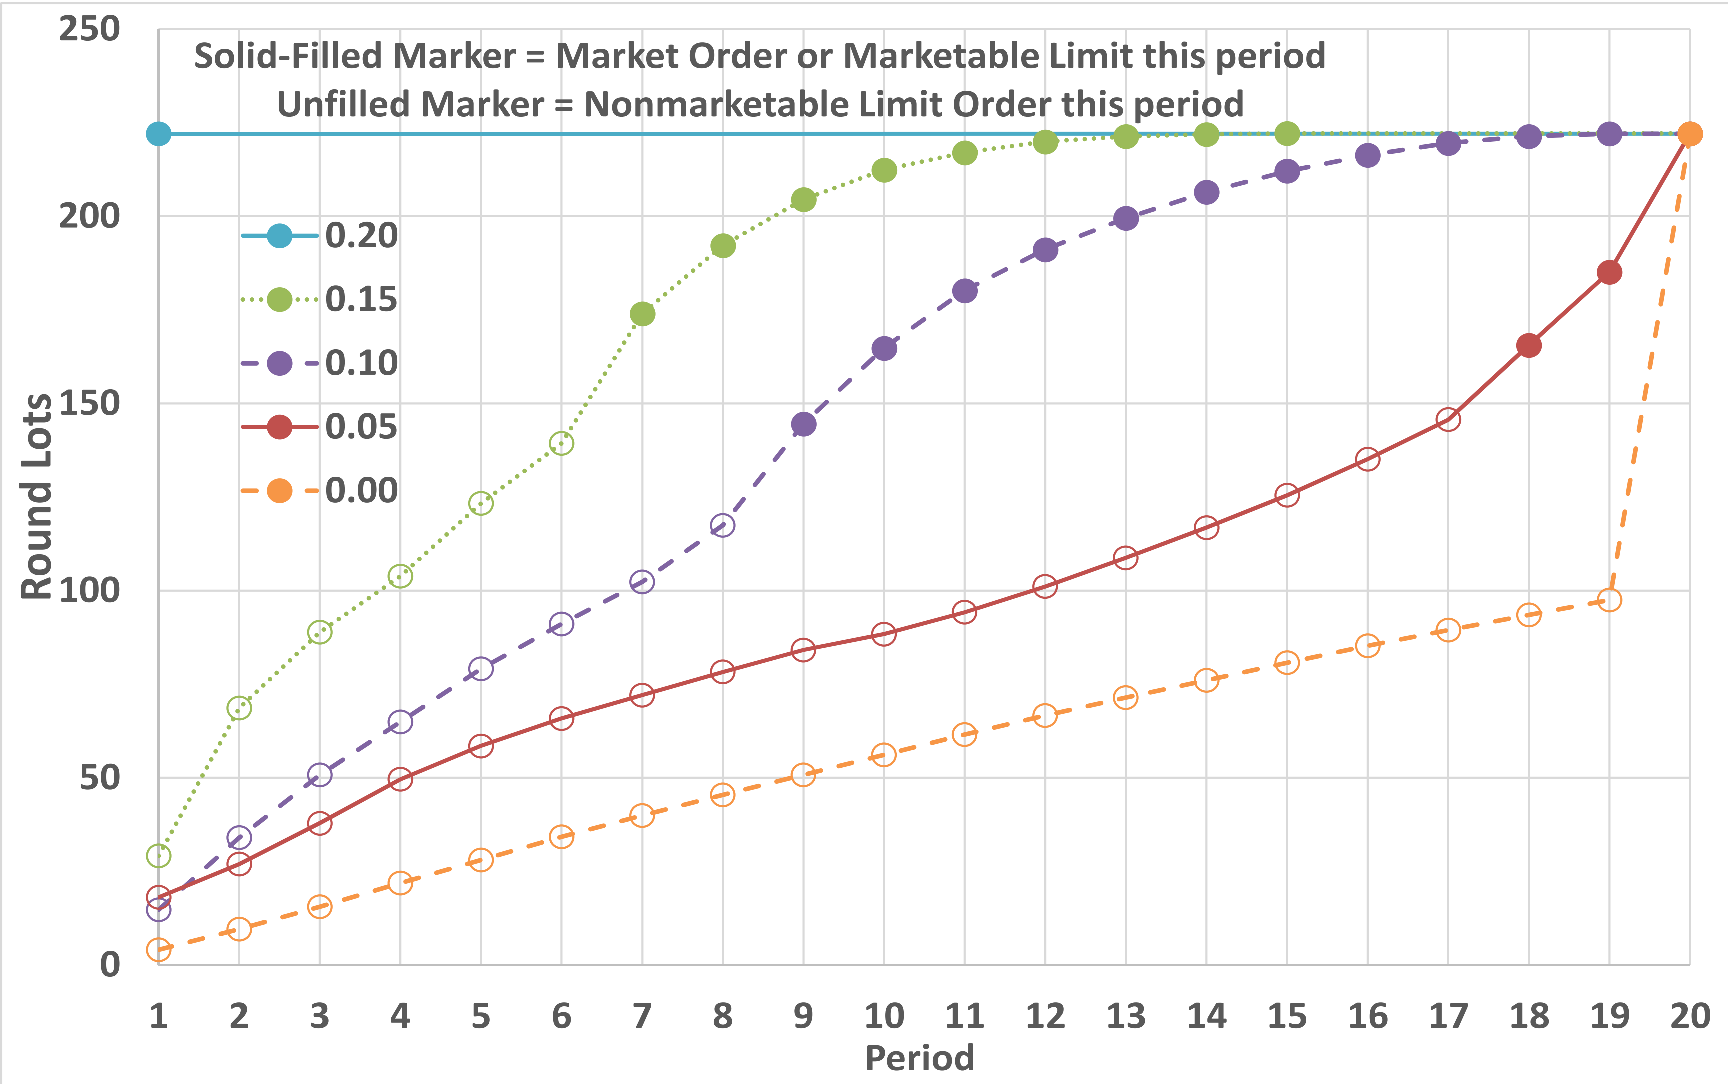

Figure 5. Average Cumulative Round Lots Bought Over Time by the Optimal Trading Strategy for Different Portfolio Manager's Degrees of Information.

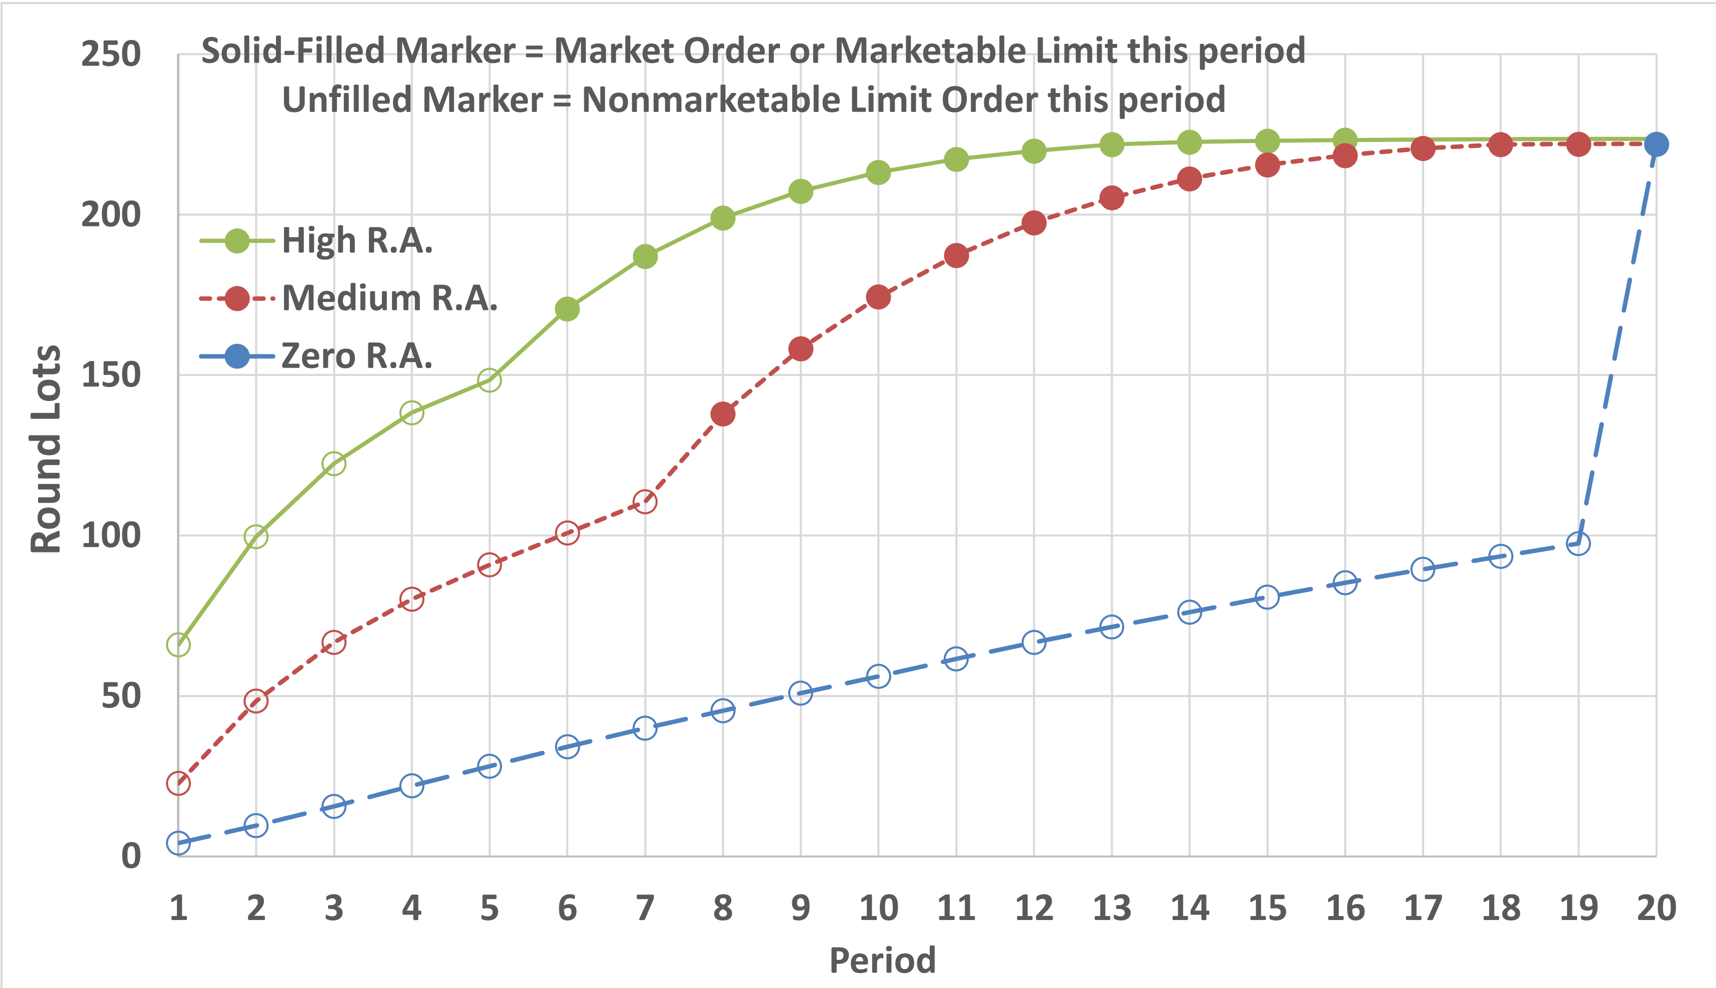

Figure 6. Average Cumulative Round Lots Bought Over Time by the Optimal Trading Strategy for Different Risk Aversions.

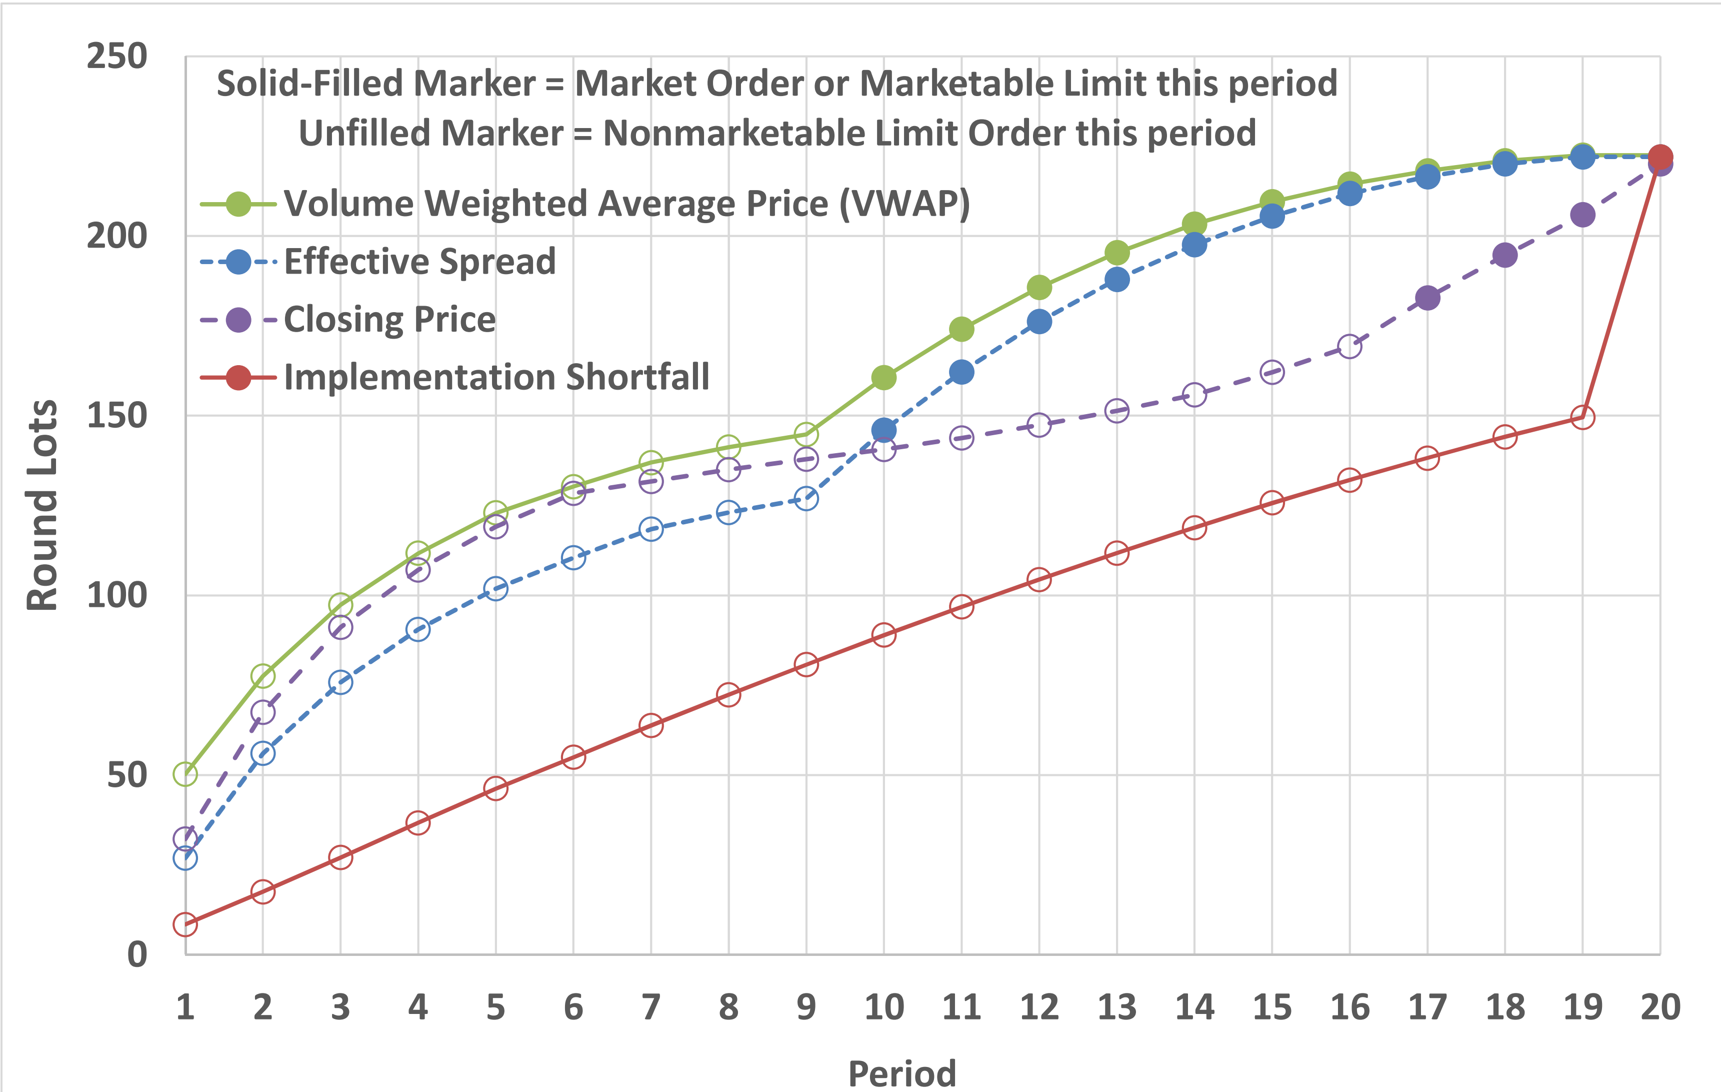

Figure 7. Average Cumulative Round Lots Bought Over Time by the Optimal Trading Strategy for Different Performance Metrics.