Optimal Dynamic Order Submission Strategies In Some Stylized Trading Problems

by

*Lawrence Harris**

*Professor of Finance and Business Economics
Marshall School of Business
University of Southern California
Los Angeles, CA 90089-1421
(213) 740-6496
LHarris@USC.edu

Draft 3.0 June 3, 1997

The author gratefully acknowledges numerous contributions made by Joel Hasbrouck. His interest and commentary have been invaluable. The author also gratefully acknowledges suggestions made by the anonymous referee, by Mark Weinstein, and by seminar participants at Hong Kong University of Science and Technology, the University of Arizona, the University of California at Riverside, the University of Iowa, the University of Southern California, the University of Utah, and Washington University.

Optimal Dynamic Order Submission Strategies In Some Stylized Trading Problems

Abstract

This study derives optimal dynamic order submission strategies for trading problems faced by three stylized traders: an uninformed liquidity trader, an informed trader and a value-motivated trader. Separate solutions are obtained for quote- and order-driven markets. The results provide practicable rules for how to trade small orders and how to manage traders. Transaction cost measurement methods based on implementation shortfall are proven to dominate other methods.

Since investors demand liquidity when they submit market orders and supply liquidity when they submit limit orders, the results improve our understanding of market liquidity. In particular, the models illustrate the role of time in the search for liquidity by characterizing the demand for and supply of immediacy.

Keywords: Limit orders, market orders, immediacy, bid/ask spreads, order submission strategies, transaction cost measurement, informed trading, value-motivated trading, dynamic programming problems.

I. INTRODUCTION

Order submission strategy is the most important determinant of execution quality that investors directly control. Traders must decide when to submit market orders and when to submit limit orders. When they submit limit orders, they must know where to place their limit prices. If their limit orders do not execute, they must know when, and how, to resubmit their orders. Investors who optimize their trading strategies will have lower transaction costs and higher portfolio returns than those who ignore their trading strategies.

Traders design order submission strategies to take advantage of the different properties of market orders and limit orders. Market orders typically produce quick executions at relatively high transaction costs. Limit orders provide lower cost executions if they execute, but they often do not execute. Limit orders also provide price-contingent executions at times when investors are unable or unwilling to monitor the market continuously.

The strategies that traders select depend on the various trading problems they solve. Traders who face early deadlines and traders who have material information that will soon become public are impatient to trade. They use market orders or aggressively priced limit orders. Conversely, value-motivated traders may be willing to wait until profitable trading opportunities arise. These traders place limit orders far from the market to represent their interests when they are unable or unwilling to continuously monitor the market.

This study derives optimal dynamic order submission strategies for several trading problems. Several problems are examined because order submission strategies depend on the different problems that traders solve. Dynamic submission strategies are examined because the search for liquidity is a sequential process: Traders often submit new orders when their limit orders do not execute. The option to resubmit thus affects the original submission decision.

The analysis considers trading problems faced by three stylized traders: an uninformed liquidity trader, an informed trader and a value-motivated trader. The liquidity trader must fill an order by a deadline. The informed trader attempts to profit from a single piece of information before it is incorporated into price. The value-motivated trader attempts to profit from on-going research into fundamental values. These three stylized problems capture most essential elements of the actual trading problems that practitioners encounter when trading small orders.[1] Although actual problems may differ from those studied here, the solutions derived in this study should provide valuable guidance to practitioners who want to derive the greatest value from their trading.

Order submission strategies also affect the supply and demand of liquidity. Traders demand liquidity when they submit market orders and they supply liquidity when they submit limit orders.[2] Efforts to understand liquidity must therefore consider how traders choose their order submission strategies.

This study contributes to our understanding of market liquidity by characterizing demands for and supplies of immediacy.[3] By analyzing dynamic submission strategies, the models presented here help characterize the relation between time and price in the search for liquidity.

The analyses do not consider how to trade large orders. Large orders are difficult to fill because they often affect prices. The models examined in this study assume that order strategy has no effect on prices. They are therefore appropriate, at best, only for small orders. Despite this limitation, the results are still useful because small orders account for a large fraction of market activity. The focus on small orders allows us to consider the relation between immediacy and the cost of liquidity, which has not been thoroughly investigated. Since size clearly affects this relation, it should be the subject of further research.

---

[1] The most notable omission from this list is the trading problem faced by traders who can substitute one security for another.

[2] A limit order instructs a broker to trade immediately at the limit price or better if possible, otherwise to hold the order until it can be executed, if ever. Since a standing limit order is an option to trade, it supplies liquidity to anyone willing to trade the other side at the limit price or better. In contrast, a market order demands liquidity by instructing a broker to trade immediately at the best price available in the market.

[3] Immediacy is the ability to trade quickly. It is one of several related dimensions of liquidity. The other main dimensions are depth (the ability to trade large size) and width (the cost of trading). Harris (1990) describes these liquidity dimensions and how they relate to each other.

The remainder of this introduction describes the three stylized trading problems. An outline of the modeling strategy is then provided and the relation of this study to prior work is described.

The formal presentation of the models appears in Section II. Section III discusses theoretical results that show how transaction cost measurement schemes affect brokers' incentives to choose order strategies that their clients prefer. Section IV provides numeric solutions to the various trading problems and discusses their properties. The text concludes in Section V with a summary and a discussion of the limitations of the analyses. Appendix A describes the methods used to evaluate the numeric integrals that appear in these problems, and Appendix B provides a detailed explanation of how price discreteness can bias the results.

A. The Three Stylized Traders

Liquidity Traders

Assume that the stylized liquidity trader is an uninformed trader who must fill an order before some deadline. Trade deadlines may arise when traders need to invest or disinvest exogenous cash flows.

Liquidity traders try to obtain the best price for their trades by carefully choosing their order submission strategies. The results presented below show that when deadlines are distant and bid/ask spreads are wide, liquidity traders initially submit limit orders. If their orders do not fill, they eventually replace them with more aggressively priced limit orders. This process continues until their trades are completed or until their deadlines arrive. If they still have not traded by their deadlines, the liquidity traders must submit market orders to assure execution.

Informed Traders

Assume that the stylized informed trader has private information about underlying value that allows him to predict future price changes. The advantage that informed traders can gain from using their information is transitory, however. These traders believe that their information will eventually become common knowledge, and that prices will change to reflect it.

Informed traders try to profitably trade on their information. In practice, trade deadlines and/or reservation price constraints often restrict their trading problems. Such restrictions may be imposed on buy-side traders by their portfolio managers or on brokers by their clients.[4] This study considers the effects of these constraints on optimal order strategy.

The models in this study assume that informed traders stop trading once they have filled a single order. In practice, traders might want to trade more intensely to profit further from their informational advantages. The decision to do so would depend on the quality of their information, their degree of risk aversion, and their access to capital.

The results presented below show that if the private information is material and if it will soon become public (i.e., if it is "hot information"), informed traders use market orders to trade quickly. Otherwise, if bid/ask spreads are wide and trading deadlines are distant, they submit limit orders to minimize their transaction costs. Those traders who face deadlines place more aggressive orders as the deadline approaches, if trading would still be profitable. Otherwise, they stop trying to trade.

Value-Motivated Traders

Assume that the stylized value-motivated trader receives a perpetual flow of private information about underlying security values. This stream of information allows value-motivated traders to forecast future values, and it allows them to model the process by which prices diverge from values.

Value-motivated traders include all traders who estimate security values on a regular basis. These traders may be investors who estimate values directly from economic fundamentals or market-makers (and other technical traders) who estimate values indirectly by observing order flows.

Value-motivated traders trade to profit from their flow of information. Since they receive continuous information about values, the model specifications assume that they can trade repeatedly.

The results reported below show that value-motivated traders demand immediacy when they believe price is far from underlying value and likely to revert quickly. Otherwise, value-

---

[4] Traders who are employed in-house by portfolio managers to manage their trades are called buy-side traders.

motivated traders offer limit orders on both sides of the market. They set their limit orders to profit from pricing errors that they anticipate might arise.

B. Overview of The Modeling Strategy

The various trading problems are modeled as dynamic programming problems. Dynamic methods are necessary because the current optimal order submission strategy depends on what future strategies will be adopted should the current strategy fail.

The model specifications involve four sets of assumptions. They describe

1. how time is modeled and when decisions are made,
2. the valuation functions that the traders maximize,
3. the set of feasible order strategies available to the trader, and
4. how the valuation function is evaluated for a given order strategy. (How execution prices and execution probabilities are determined.)

This section provides an overview of these assumptions and of how the models are solved and analyzed. The formal presentation of the models begins in Section II.

Time Intervals

The models assume that traders submit and/or cancel their orders only at discrete time intervals. In practice, traders make order decisions at discrete intervals because they cannot continuously monitor and adjust their orders, because their time is valuable, and because canceling and resubmitting orders is costly.

The inability to continuously adjust orders gives limit orders option-like properties. Copeland and Galai (1981) show that limit orders are like options because they give traders an option to trade at a fixed price. The longer the order is expected to stand, the more valuable the option will be. In volatile markets, traders place their limit orders far from the current price to reduce these option values.

Assume that traders submit a single order (if any) at the beginning of a time interval. If an order does not execute by the end of the interval, the trader may cancel it and resubmit a new order in the next interval. In practice, traders may submit multiple orders in the same security to obtain the best possible average price for a large transaction. This analysis does not consider such strategies because the assumptions made in this study are appropriate only for small orders.

This analysis assumes that the time interval is of fixed length. In practice, cancellation and resubmission decisions probably also depend on unexpected events.[5] Depending on the volatility process assumed below (stochastic or homoskedastic), the fixed interval length may be interpreted as a fixed chronological time interval or a fixed event-time interval.

The models do not specify the chronological (or event-time) length of an interval. Instead, the interval length is implied from the volatility assumed for the price process within the interval. For the same assumed degree of volatility, the model may represent a trading problem involving a volatile security examined at short intervals or a stable security examined at long intervals. An interval therefore may be a day, an hour or a minute. In actual practice, the length of an interval is the time over which the trader is unwilling or unable to revise his order.

The models also allow for (but do not require) nontrading periods between intervals. Such periods might be overnight periods or lunch time recesses.

Valuation Functions

The valuation function specifies how traders value order outcomes. This study examines several different valuation functions because trader objectives depend on their trading problems. Traders who are precommitted to trading typically want to maximize price when selling and minimize price when buying. Informed traders and value-motivated traders, however, generally want to maximize their trading profits. These objectives differ because informed and value-motivated traders do not have to trade.

Brokers' objectives may differ from their client's objectives depending on how their clients measure and evaluate their performance. Since different performance evaluation methods may cause brokers to adopt different strategies, this study considers valuation functions based on several different transaction cost evaluation schemes.

---

[5] Evidence in favor of the mixture of distribution hypothesis, such as that presented in Harris (1986a,1987), suggests that many trading decisions are made in event-time.

This study examines valuation functions based only on expected trading profits and expected transaction costs. The traders are therefore assumed to be risk neutral. Risk neutrality does not seriously reduce the usefulness of the results even if traders are risk averse. Expected profit maximization will generate approximately the same strategies as expected utility maximization if order size is small relative to wealth and if the trading problem is frequently repeated.[6]

Feasible Strategies

The set of feasible strategies specifies the actions that traders may take at the beginning of each interval. This study assumes that at each interval, a trader can submit a market order, a limit order with a discrete limit price, or do nothing.

Sometimes the trading problem restricts the feasible set. For example, precommitted liquidity traders must submit market orders at their trading deadlines, and brokers cannot submit orders that would violate reservation price constraints.

Valuation Function Evaluation

The final set of assumptions specifies how the traders' expect their orders will perform and how they expect prices to evolve over time. These assumptions are used to compute

1. the prices at which orders execute,
2. the probabilities that they execute, and
3. the distribution of future prices should orders fail to execute.

These quantities are then used to evaluate the valuation function for a given order strategy. Since the three stylized traders have different private information about security values, different sets of assumptions characterize the traders' expectations.

Trader expectations about how orders execute depend on the market structure into which they submit their orders. For example, in pure quote-driven dealer markets, small orders typically execute at the best opposing dealer quote, regardless of whether they are market or limit

---

[6] Expected utility objective functions can be analyzed with the methods used in this paper, but at substantially greater numeric cost. In the risk neutral models presented below, risk aversion can be represented approximately by assigning fixed costs for submitting limit orders that produce uncertain execution prices.

orders.[7] In public auction markets, market orders sometimes experience price improvement, and limit orders execute with probabilities that depend on where traders place their limit prices relative to the standing bid and ask quotes.[8] The models introduced in this study allow these differences to be specified.

Trader expectations about order execution also depend on security characteristics such as volatility, spread, and liquidity. For example, wide spreads make market orders expensive, and high volatility facilitates limit order executions. Again, the models allow these characteristics to be specified.

Solution and Analytic Methods

The dynamic programs specified in this study are solved using numeric methods. Numeric methods are necessary because the integrals in the valuation functions that represent trader expectations are analytically intractable. After evaluating these integrals, the solutions are obtained using standard dynamic programming methods. The need to evaluate numeric integrals of numeric functions, however, complicates the analysis.

The dynamic programs include parameters that specify characteristics of the valuation function, of the feasible set, and of trader expectations about prices and order execution processes. The most important of these parameters characterize price volatility, the bid/ask

---

[7] A pure quote-driven market is a market in which dealers supply all the liquidity. In such markets, public traders typically buy from dealers at the lowest ask price offered by any dealer, and sell to dealers at the highest bid price offered by any dealer. Although dealers may represent a public order in their quote, they rarely do so.

[8] In a public auction market, an exchange matches public orders with each other according to a set of precedence rules. Precedence usually is based first on price and then on time of arrival. In a continuous public auction market, trades are usually arranged when a market buy (or sell) order is matched to the standing sell (or buy) limit order that has the greatest precedence. A public limit order can therefore supply liquidity directly to a public market order. A limit order that is aggressively priced (high price if a buy order, low price if a sell order) has high precedence and therefore will more likely execute quickly than will a less aggressively priced order. Dealers may also participate in these public auction markets. Market orders experience price improvement if a trader—typically a dealer—is willing to fill the order at a price better than the quote.

spread, the minimum price variation (tick size), information dissemination rates, and various aspects of execution difficulty. Comparative static effects for each parameter are identified by solving the models for various parameter values.

C. Relation to Other Studies

The present study is related to three themes in microstructure research. The first theme considers how transaction costs should be measured to evaluate and control trader performance. Berkowitz, Logue and Noser (1988), Beebower (1989) and Perold (1988) propose several transaction costs measurement methods. Beebower (1989) and Harris (1990a) informally evaluate their various properties. The present study provides formal theoretical results that show which measurement methods best solve the agency problems associated with trading through intermediaries.

The second theme examines empirical evidence on order submission strategies. With the increasing availability of order flow data, several studies attempt to describe the ex post performance of particular order submission strategies. Bronfman (1992), Angel (1995) and Harris and Hasbrouck (1996) examine the tradeoffs between NYSE market and limit orders. The strengths and limitations of these papers lie in their reliance on actual orders. When the history of an order is fully described over time, its final disposition can be determined precisely. On the other hand, these studies have little or no information about the trader’s overall objective, the market conditions that led the trader to select a particular order, or how the order fit in with the trader’s dynamic strategy.

Without individual order data, inferences may be drawn from the patterns of reported trades and quotes. In the CAC system used by the Paris Bourse, for example, quotes derive almost exclusively from public limit orders and transactions reflect the crossing of market (or marketable limit orders) with these quotes. Biais, Hillion and Spatt (1995) examine order arrival characteristics in this market. Hamao and Hasbrouck (1992) likewise consider the Tokyo Stock Exchange, which is similarly organized. The market data underlying these studies typically permits computation of execution costs for market orders. Even for market mechanisms that do not permit direct inferences about orders from reported trade and quote data, it may be possible to investigate simulated strategies. Handa and Schwartz (1993) examine the performance of hypothetical limit order strategies given actual New York Stock Exchange price paths.

The final theme considers how traders form their order submission strategies. Several theoretical analyses examine models of trader order submission decisions. These models are typically single-period (or at least single decision-point) analyses in which the trader may place a market order or a limit order for a given quantity. In Cohen, Maier, Schwartz and Whitcomb (1981), this choice (under homogeneous information) leads to a non-trivial spread between the highest limit buy price and the lowest limit sell price. At the equilibrium spread, the agent’s cost (in expected utility) of submitting a market order and paying the spread is exactly equal to the cost of a limit order — a cost arising from the possibility of non-execution. Kumar and Seppi (1993) model order submission with informed and uninformed traders. Chakravarty and Holden (1995) examine one-period order submission strategies for an informed trader. The present study examines trader problems and strategies that are considerably more general than these analyses, but it does not attempt to characterize equilibria.

Angel (1995) presents a model in which a trader who must trade chooses between a market order and a limit order. If the limit order fails to execute, the trader resubmits the order as a market order. The market environment (other traders' market and limit orders) are assumed to arise from random Poisson arrivals at random prices. The present analysis is less detailed in its description of the environment in that the characteristics of other traders' orders are subsumed in a reduced-form model. It is more general, however, in its consideration of multiperiod dynamics, allowance for order revision and range of other trading problems.

The study most closely related to this one is Bertsimas and Lo (1996). They derive solutions to a dynamic programming problem in which a large trader spits orders in an attempt to minimize transaction cost. The present study does not consider the effect of large orders on price. It does, however, examine a wider variety of trading problems than does Bertsimas and Lo. In addition, it examines the incentive compatibility of various used to evaluate brokers.

II. THE FORMAL MODELS

This section provides formal specifications of the various dynamic programming models. The presentation starts with features common to all three stylized problems and then proceeds to features that differ across the problems. For the reader's convenience, Tables 1 and 2 present annotated lists of all notation and assumptions introduced in this section.

A. TIME INTERVALS

The following assumptions characterize time in the dynamic trading problems:

1. Time is divided into a series of discrete intervals.
2. Traders submit an order (if any) only at the beginning of an interval.
3. Traders may submit orders only if the market is open.
4. Prices and values change within intervals and trades take place within intervals.
5. Between intervals, prices and values may change but trades do not take place.
6. All unfilled orders expire at the end of each interval.

The following conventions are used to refer to time: Intervals are indexed by the subscript t. The index τ refers to time within an interval. For notational convenience, τ ranges from 0 to 1 so that $\tau = 0$ indicates the beginning of an interval and $\tau = 1$ indicates the end. Variables subscripted by t but not indexed by τ refer to beginning of interval values. Likewise, variables subscribed by $t+$ but not indexed by τ refer to end of interval values. The period between $t+$ and $t+1$ is called the nontrading interval.

B. VALUES, PRICES AND INFORMATION

This subsection and the next describe the processes by which traders expect prices to evolve and orders to execute. The assumptions that characterize these processes are important because they determine the value of alternative order submission strategies.

This subsection specifies how traders expect prices to evolve. Separate assumptions are made for each stylized trader. To avoid repetition, assumptions common to each trader are presented together. For the reader's convenience, Table 2 summarizes the three sets of assumptions.

Trader price expectations are developed by specifying three processes. These processes describe the evolution of values, pricing errors, and pricing error estimates.

The first process describes how fundamental security values change. Fundamental security value is the (unobserved) theoretical market value of the security that all traders would agree upon if all information about the security were common knowledge, if all traders were superior analysts, and if all traders traded rationally. The difference between a security's market price and fundamental value is called the pricing error.

The second process describes how pricing errors evolve. This process is not specified for the uninformed traders. It is specified only in schematic form for the informed traders, and it is formally specified for the value-motivated traders.

The final process describes what traders know about the current pricing error, and what they expect about future pricing errors. Trader pricing error estimates depend on the information they have about values, prices and about other traders. The means by which traders obtain their estimates from this information are not formally modeled in this study because doing so would render the resulting models intractable. Instead, the process is characterized by reasonable assumptions about the resulting pricing error estimates.

Values

Assume that each stylized trader believes that within an interval, underlying security value follows Brownian motion with zero drift and that volatility varies by interval.[9] Let $V_t(\tau)$ denote value at time τ within interval t, let

---

[9] Brownian motion is assumed because the value diffusion process must be symmetric within intervals to obtain tractable solutions. Zero intra-interval drift is assumed to keep the computations manageable. Allowing drift would increase the dimension of the numeric integrals necessary to solve the dynamic programming problems by one. The zero drift assumption should not change the character of the problem significantly since unconditional expected changes in value are small within typical trading intervals.

$$\begin{aligned} \varepsilon_t &= V_t(1) - V_t(0) \\ &= V_{t+} - V_t \end{aligned} \tag{1}$$

denote the interval change in value, and let $\sigma^2_{V,t}$ denote the constant volatility of the Brownian motion within interval *t*. These assumptions imply that the traders believe the interval value change ε_t is conditionally normally distributed with mean 0 and variance $\sigma^2_{V,t}$.

Assume that traders believe that the variance $\sigma^2_{V,t}$ follows an i.i.d. Inverted Gamma process.[10] The mixing of the conditional normal distribution with this volatility distribution produces an unconditional Student-*t* distribution for value changes. Trader expectations thus reflect the well-known fat-tailed properties of price changes. Let ν denote the Inverted Gamma (Student-*t*) degrees of freedom parameter. For small ν, the Inverted Gamma distribution is highly right skewed so that the value process will seem as though it includes a jump process. Modeling extreme event probabilities is important because limit order option values (and therefore optimal limit order submission strategies) are sensitive to extreme events.

Finally, assume that traders believe values may also change between intervals. Let

$$\begin{aligned} \theta_t &= V_{t+1}(0) - V_t(1) \\ &= V_{t+1} - V_{t+} \end{aligned} \tag{2}$$

denote the nontrading interval change in value following interval *t*. Assume that traders believe that θ_t is independently distributed with mean μ that reflects an expected change in value. For most analyses presented in this study, the shape and dispersion of the nontrading interval value change distribution do not matter because the traders are assumed to be risk neutral, and because their objective functions are all linear in trading profits or transaction costs. Solutions to trading problems that include a reservation price constraint do depend on the inter-interval volatility process because the inter-interval volatility affects the probability that the constraint will be binding in the future. For these problems, assume that that traders believe that θ_t is independently and identically normally distributed with mean μ and conditional variance

[10] Autoregressive heteroskedasticity could be modeled by adding another state variable to the dynamic programming problems. This addition, however, would make coding and computing solutions to the problem prohibitively expensive.

proportional to $\sigma^2_{V,t}$ (so that there is a common volatility factor), and assume that θ_t and ε_t are independent.

Prices and Information

Assume that traders believe that the midspread price *P* (the average of the bid and ask quotes) is the sum of the underlying security value *V* plus a pricing error *e*:

$$P_t(\tau) = V_t(\tau) + e_t(\tau) \tag{3}$$

The bid *B* and the ask *A* are respectively one-half of the spread *S* below and above the midspread price *P*:

$$\begin{aligned} B_t(\tau) &= P_t(\tau) - \tfrac{1}{2}S \\ A_t(\tau) &= P_t(\tau) + \tfrac{1}{2}S \end{aligned} \tag{4}$$

The three stylized traders differ in the information they have about the pricing errors.

Assume that uninformed liquidity traders have no information about pricing the errors and security values. Further assume that they have no model that would allow them to extract information about the pricing error process from the history of prices. With such poor information, prices appear to follow a random walk to them. Their best estimate of the pricing error will be zero:

$$\hat{e}_t(\tau) = 0\ . \tag{5}$$

Informed traders have information that allows them to construct a one-time estimate of value, and from that estimate, an estimate of the current pricing error. Assume that the error in their pricing error estimate, $\psi_t = e_t - \hat{e}_t$, has mean zero and is independent of all other variables.

Informed traders believe that the information upon which their value estimate is based will eventually become public and that prices will adjust accordingly. As time passes, their estimate of the pricing error must therefore decline towards zero. Assume that informed traders expect that their estimate of the pricing error will decay at a constant exponential rate.

The actual rate of decay in their pricing error estimate may be different from the expected rate of decay. Even though informed traders receive no new private information about values, they may be able to make some inference about values from subsequent prices. For example, suppose that based on their information, informed traders believe price is 30 percent below fundamental value and that 10 percent of the difference (3 percent) will be realized within one day. Suppose further that they believe that the standard deviation of value innovations is two percent per day. If price rises by 15 percent in one day, these traders will be surprised: They expect price to rise by only 3 percent as other traders learn their information, plus or minus about two percent for unexpected changes. Upon seeing a 15 percent rise, these traders will more likely attribute it to an earlier than expected reaction to their information than to an extraordinary 12 percent rise in value. Their estimate of the pricing error on the next day will be closer to 15 percent than to the 27 percent that they earlier expected. In practice, their estimated pricing error will be a function of all price changes observed since they obtained their information.

To keep solutions tractable, this learning process is not modeled. By assuming that informed traders expect that their estimate of the pricing error will decay at a constant exponential rate independent of future price changes, the model implicitly assumes that informed traders are myopic about the effect of future price changes on their future estimates of value.[11] This assumption does not, however, preclude the traders from subsequently updating their estimates after prices have changed.

To keep the dynamic programs computationally manageable, assume that informed traders expect that the decay in their estimated pricing error takes place only between intervals so

---

[11] The bias in the optimal trading strategy that results from this myopic assumption is not likely to be large. The bias depends on the implied bias in the expected valuation function. Assume that the myopic estimate is the unconditional mean of the proper update. The bias in the expected valuation function can then be approximated using a simple second order Taylor series expansion around the myopic estimate. The resulting approximate bias is the product of the variance of the myopic approximation error and the second derivative of the valuation function about the myopic expected pricing error. When the expected pricing error is small, the bias in the optimal strategy is likely to be small because the myopic approximation error variance will be small. When the expected pricing error is large, the bias in the optimal strategy is likely to be small because the value function will be a nearly linear function of the expected pricing error and because the optimal strategy will probably be a market order in any event (if the expected decay rate is significant.)

that the expected price change within the interval is zero.[12] Stated formally, at the beginning of interval t, informed traders make an estimate (assumed to be given in this problem) $\hat{e}_t$ of the pricing error e_t based on their initial estimate of value and perhaps also on changes in price observed since then. Their expectations for the evolution of their pricing error estimate are described by

$$\mathrm{E}\left(\hat{e}_t(\tau)\right) = \hat{e}_t \tag{6}$$

within interval t, and

$$\mathrm{E}\left(\hat{e}_{t+1}\right) = \phi \hat{e}_t \tag{7}$$

between intervals t and t+1 where $(1-\phi)$ is the exponential rate of decay. This assumption is unattractive because the incentives to trade quickly are not properly represented when the expected rate of decay is high. This deficiency can be addressed by shortening the interval length (by reducing variances) so that the expected end-of-interval decay comes sooner.[13]

Value-motivated traders continuously obtain information that they use to estimate current security values and pricing errors. Let the error in their pricing error estimate, $\psi_t = e_t - \hat{e}_t$, have mean zero and be independent of all other variables.

Since value-motivated traders estimate pricing errors continuously, they should know the process by which pricing errors evolve. Assume that value-motivated traders believe that the pricing error within the interval follows Brownian motion. Innovations in the pricing error process may be due to the random arrival of liquidity traders or to innovations in the value process that are not public knowledge. Like the informed traders, value-motivated traders also believe that their information will eventually become public and that prices will adjust accordingly. For reasons identical to those discussed above, assume that value-motivated traders

---

[12] Zero drift in the price diffusion process within the interval greatly simplifies the derivation of first passage probabilities.

[13] Mathematical results in Shepp (1979) would allow the expected decay to be modeled as a linear drift in prices within the interval. Although this specification would more realistically represent the incentives to trade quickly within an interval, it was not undertaken because its implementation would increase the dimension of the numeric integrals by one. The resulting complexity would have made the problems too computationally burdensome to solve.

expect that the pricing error will decay at a constant exponential rate, and that the decay takes place only between intervals.

The value-motivated traders' model for the pricing process is formally described as

$$e_t(\tau) = e_t + n_t(\tau) \tag{8}$$

within interval *t*, and

$$e_{t+1} = \phi e_{t+} \tag{9}$$

between intervals *t* and *t*+1, where $e_{t+} = e_t(1) = e_t + n_t(1)$. The pricing error innovation $n_t(\tau)$ follows zero drift Brownian motion with innovation variance $\sigma^2_{N,t}$. The covariance between the value innovations and the pricing error innovations is given by $\sigma_{VN,t}$. So that time in the pricing error process remains synchronized with time in the value process, assume the two variances and the covariance are all proportional to the same volatility factor. These assumptions imply that the value-motivated traders believe price changes measured across intervals follow an ARMA(1,1) process.

At the beginning of each interval *t*, value motivated traders estimate the current pricing error $\hat{e}_t$. Using their assumed model for the pricing error process, their implied expectation of $\hat{e}_{t+1}$ at the beginning of interval *t* is

$$\mathrm{E}(\hat{e}_{t+1}) = \phi \hat{e}_t . \tag{10}$$

C. Execution Mechanism Expectations

This subsection describes the models used to represent how traders expect their limit and market orders will execute. This exposition and all subsequent discussions in this study assume that the traders want to sell. The analysis for buy orders would be symmetric.

Market orders

Traders assume that market orders execute immediately when they are submitted at the beginning of the interval. The execution price depends on the market structure, the bid/ask spread and on cross-sectional characteristics of the security. If a small market sell order is sent to a pure quote-driven dealer market, the execution price will typically be the dealer bid. If the order is sent to an auction market, it may sometimes execute at a price inside the spread. It may occasionally even execute at the ask if the order is stopped. The difference between the actual trade price and the bid is called the price improvement.

Solutions to the dynamic trading problems analyzed in this study require only a model for the average market order execution price. Assume that the worst possible execution price for a sell market order is the beginning of interval bid price, and assume that traders expect that the average price improvement is some constant fraction π of one-half the bid/ask spread. In a pure quote-driven dealer market, traders expect that $\pi = 0$. In an auction market, traders will expect π to be between 0 and 1.[14] To summarize, the expected price for market orders submitted at the beginning of interval t is

$$\begin{aligned} EM_t &= B_t + \pi \tfrac{1}{2} S \\ &= P_t - (1-\pi) \tfrac{1}{2} S \end{aligned} \tag{11}$$

This analysis treats marketable limit orders as market orders. A marketable limit order is an order that can be immediately executed against the opposite side quote at the time of submission. For example, a sell limit order is a marketable order if its limit price is equal to or lower than the best bid in the market (and the quantity attached to the best bid is sufficient to satisfy the limit order size).

Limit orders

The limit order execution mechanism also depends on the market structure. For example, in pure quote-driven dealer markets, a sell limit order typically executes only when the highest dealer bid rises to the limit sell price.[15] In auction markets, a sell limit order executes only when it represents the best offer and somebody is willing to take the other side of the trade.

---

[14] In specialist auction markets when the spread is one-eighth, sell market orders occasionally execute at the ask. Although the implied price improvement parameter for such executions is greater than one, the average price improvement cannot be greater than one if the specialist does not favor sellers over buyers and if the specialist does not intentionally lose money on market orders.

[15] This condition does not guarantee execution in some dealer markets. If the dealer who holds the order is not required to execute it when his bid is not the best bid, limit order execution may be even more difficult.

This study analyzes two limit order execution mechanisms. The first mechanism, which is a special case of the second, is most useful for describing limit order execution in pure quote-driven dealer markets. The second mechanism provides a more realistic execution model for auction markets and negotiated dealer markets. Most results in the study are based on the second mechanism.

For the first execution mechanism, assume that a sell limit order will execute if, at any time within the interval, the midspread price exceeds the limit order price by some constant. The constant determines the difficulty of execution. Its value depends on the market structure, the bid/ask spread and on cross-sectional characteristics of the security. In a pure quote-driven dealer market in which a limit sell order only executes when the best dealer bid rises to the limit sell price, the constant is equal to one-half of the bid/ask spread. In very active auction markets, in which a steady stream of market orders quickly executes standing limit orders, the constant is minus one-half of the bid/ask spread. This extreme assumption implies immediate execution for any sell order placed at the offer at the beginning of the interval. Increasing the constant will produce more realistic results. By varying the constant between minus one-half and one-half of the spread, different degrees of execution difficulty can be represented. In all cases, the execution price is equal to the limit order price. To summarize, sell limit orders execute under the first execution mechanism if

$$L_t < P_t(\tau) - \delta \text{ at any } \tau \text{ in interval } t. \tag{12}$$

where L_t is the limit price of a limit order submitted at the beginning of interval t, and δ is the degree of execution difficulty parameter. Let $\Pr(Exec|\delta)$ represent the probability that the order executes given execution difficulty δ. For a given limit price L_t, $\Pr(Exec|\delta)$ decreases with an increase in δ. Appendix A provides a formula to evaluate this probability.

The second execution mechanism provides a more realistic model of execution. Under the first mechanism, a sell limit order executes **with certainty**, if at any time within the interval, the midspread price exceeds the limit order price by the execution difficulty constant.[16] Under

[16] In particular, note that if the limit price is set at $P_t(0) - \delta$, the limit order will execute immediately because the probability that Brownian motion rises above its starting point within any finite interval is one.

the second mechanism, assume that execution is not certain when this condition is met. Assume instead that the sell limit order will only execute **with some probability** when the condition is met.

Call this probability the fill probability and assume that it is an increasing function of the execution difficulty parameter δ. For δ greater than or equal to one-half of the spread, execution should be certain when the condition is met: Satisfaction of the condition implies that the bid will have crossed the sell limit price at some point. For δ less than minus one-half of the spread, no execution should occur if the execution condition is only just met but not exceeded. In this case, the sell limit price will never have been lower than the ask.

Assuming that the probability of execution is the weighted sum of execution probabilities computed under the first mechanism completes the definition of the second execution mechanism. The sum is computed over all possible δ and the weights are given by the density of the fill probability function.

The second execution mechanism can be described formally as follows: Let $F(\delta)$ be the fill probability function and let $f(\delta)$ be the associated density function. Assume that F takes the following form:

$$F(\delta)=\begin{cases}0 & \text{for } \delta<-\tfrac{1}{2}S \\ G(\delta) & \text{for } -\tfrac{1}{2}S\le\delta\le\tfrac{1}{2}S \\ 1 & \text{for } \delta>\tfrac{1}{2}S\end{cases} \tag{13}$$

where $G(\delta)$ is a distribution function rising from 0 to 1 over the interval $\left[-\tfrac{1}{2}S,\tfrac{1}{2}S\right]$. The probability that an order is executed is

$$\Pr(Exec)=\int_{-\infty}^{\infty}\Pr(Exec|\delta)\,f(\delta)\,d\delta\,. \tag{14}$$

Appendix A provides a formula to evaluate this probability. The first execution mechanism is a special case of this second mechanism when

$$G(\delta)=\begin{cases}0 & \text{for } \delta\le\delta^{*} \\ 1 & \text{for } \delta>\delta^{*}\end{cases} \tag{15}$$

where δ^{*} is the given specific execution difficulty parameter.

This analysis uses a Beta distribution to represent $G(\delta)$. The standard Beta distribution is shifted and scaled so that it ranges over the interval $\left[-\frac{1}{2}S, \frac{1}{2}S\right]$. The two degrees of freedom parameters in the Beta distribution allow the specification of a wide variety of shapes. They also determine the mean and standard deviation of the distribution. When both degrees of freedom are equal to one, the resulting Beta distribution is the Uniform distribution.

D. VALUATION FUNCTIONS

This subsection describes the valuation functions that the various traders maximize. For notational convenience, consider only sales of a single share or contract.

Liquidity Traders

Assume that liquidity traders are risk neutral so that they simply want to maximize the expected net sales price. (The net sales price is the sales price less any costs incurred through order submission and execution.)

Their brokers, however, may have different objectives. Since customers usually cannot monitor their brokers' efforts, brokers may shirk. To prevent this problem, customers often evaluate brokers by comparing their trade prices to some price benchmark. A risk neutral broker's objective may therefore be to minimize measured transaction cost (the difference between the trade price and the benchmark price) rather than to maximize the sales price.

This study examines three price benchmarks commonly used to measure transaction costs: the beginning-of-interval bid price B_t taken from the interval during which the trade takes place, the corresponding end-of-interval bid price B_{t_+}, and a prespecified price, P_o. (When P_o is the midspread price prevailing at the time of order submission, the transaction cost measurement method is called the implementation shortfall method.) Bid price benchmarks are examined because the resulting valuation function is easily interpreted: Any improvement in price from using a limit order is measured relative to the worst price at which a market order would have executed. Benchmarks based on ask and midspread prices produce identical optimal order strategies to those based on the bid price because the three price processes differ from each only by constants.[17]

---

[17] This obvious result is formally proven in Theorem 1 of the next section.

The formal description of the valuation function consists of two parts. The first part describes the dynamic objective function while the second part describes the terminal value of the valuation function. The dynamic part is given by

$$W_t(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{if a limit order is submitted but does not fill} \\ L_t - BMP_t - K_L - C_L & \text{if a limit order is submitted and filled} \\ M_t - BMP_t - K_M - C_M & \text{if a market order is submitted and filled} \\ W_{t+1}(P_{t+1}) & \text{if no order is submitted} \end{bmatrix} \tag{16}$$

where $W_t(P_t)$ is the value of the dynamic program at the beginning of interval t given the midspread price P_t , L_t is the limit price of a limit order submitted at the beginning of interval t, M_t is the execution price for a market order submitted at interval t, BMP_t is a benchmark price, and K_L , K_M , C_L and C_M are the costs of submitting and the costs of executing limit and market orders. When the benchmark price is zero, the valuation function represents the sales price maximization objective. Otherwise the valuation function is a simple transform of the transaction cost minimization objective: Maximize the negative of the estimated transaction cost.

The liquidity traders' trade deadline implies a terminal value for the valuation function. Since the traders must use market orders to complete unfinished trades at the deadline T, the terminal value for the dynamic program is given by

$$W_T(P_T) = \text{E}\left(M_T - BMP_T - K_M - C_M\right). \tag{17}$$

The valuation function depends on two state variables, time and price. The deadline ensures that it depends on time. The valuation function depends on the current price because the expected prices at which market and limit orders execute depend on the current price. However, when the benchmark price is a beginning-of-interval price or a post-interval price, the difference between the expected trade price and the expected benchmark does not depend on the current price. For these benchmarks, the valuation function depends only on time. When the price benchmark is a constant prespecified price, the valuation function depends on price, but the solution to the problem fortunately does not: Corollary 1 of Theorem 1, presented below in Section III, shows that the solution to the prespecified price benchmark dynamic program is exactly equivalent to the solution to a related program having a valuation function that depends only on time. The reduction in the state space is accomplished by manipulating the problem so that all prices are expressed relative to the current midspread price.

The maximization is conducted over the set of feasible order submission strategies (described below). The expectation is computed over all possible outcomes, whose generating processes are described in the previous section.

Informed Traders

The informed trader's objective is to obtain the maximum price for the sale subject to the reservation price constraint. If the reservation price constraint is set at the trader's estimate of value, this objective maximizes the value of the portfolio. As in the liquidity trader's problem, brokers who trade to minimize estimated transaction costs may choose different strategies than are optimal for sales price maximization.

This study considers valuation functions based on several benchmark prices. When the benchmark price is zero, the objective is to simply maximize the sales price. When the benchmark price is the future value of the security, V_{t+s}, the objective is to maximize the expected value added from trading. When the benchmark price is the price prevailing at the time of order submission, the objective is to minimize the implementation shortfall measure of transaction costs. Finally, when the benchmark price is B_t, the beginning-of-interval bid price or B_{t+}, the end-of-interval bid price, the objective is to minimize transaction costs measured relative to these prices.

For a given benchmark price, the informed trader problem valuation function is similar to the liquidity trader problem valuation function.[18] The expression for the dynamic objective function is identical to (16). Expressions for the valuation function terminal values differ, however, because the informed traders can choose not to trade. For the sales price maximization objective, the terminal value of the valuation function is

$$W_{T+1}\left(P_{T+1}\right) = \mathrm{E}V_{T+1} \tag{18}$$

[18] Although the valuation functions have similar expressions, the evaluation of their expected values differ because the two types of traders have different information.

since the trader believes that had the security been bought, it would have been worth V_{T+1} at the end of the deadline interval. For the value-added maximization objective, the terminal value is simply zero since no value is added if the trader does not trade:

$$W_{T+1}\left(P_{T+1}\right) = 0\,. \tag{19}$$

For the implementation shortfall minimization objective, the terminal value is the lost profit from the trade not completed:

$$W_{T+1}\left(P_{T+1}\right) = P_{T+1} - P_o \tag{20}$$

where P_o is the mid-quote price at the time of order submission. For the transaction cost measures based on beginning- and end-of-interval prices, the terminal value is zero as in (19) because transaction costs generally are not computed for trades not completed (except for the implementation shortfall method).

Value-motivated Traders

Value-motivated traders trade to maximize the values of their portfolios. This objective differs from seeking the highest price for each trade because value-motivated traders may return to the market after each trade is completed. Value-motivated traders may therefore trade more aggressively (accept less desirable prices) to increase their rate of profitable trades.

The value-motivated traders' dynamic valuation function is

$$W\left(\hat{e}_t\right) = \beta \mathrm{E} W\left(\hat{e}_{t+1}\right) + \text{Max E} \begin{bmatrix} -K_L & \text{if a limit order does not fill} \\ L_t - V_{t+} - K_L - C_L & \text{if a limit order fills} \\ M_t - V_{t+} - K_M - C_M & \text{if a market order fills} \\ 0 & \text{if no order is submitted} \end{bmatrix} \tag{21}$$

where β is an inter-temporal discount factor. The discount factor is necessary because the problem has an infinite horizon. The valuation function W represents the present discounted value of trading an infinite series of sell orders. It depends on $\hat{e}_t$, the value-motivated trader's current estimate of the pricing error. Since the value function does not appear on any of the lines over which the maximum is computed, it has no effect on the optimum strategy.[19]

E. THE FEASIBLE SET

Each maximization is conducted over a set of feasible order submission strategies. This set includes a market order, limit orders set at various discrete prices, and the possibility of no order submission. The discrete limit price set is given by

$$\left\{L_t = A_t \pm id \text{ for } i = 0,1,\dots,\infty\right\} \tag{22}$$

where A_t is the current ask price and d is the discrete tick.[20]

When the seller is subject to a reservation price constraint, the constraint imposes further restrictions on the set of feasible strategies. Limit order prices cannot be less than the reservation price ($L_t \geq P^R$) and no market order may be submitted if the worst execution price (B_t) could be less than the reservation price.

---

[19] Since the order strategy does not depend on the value function, the value function for the combined problem of setting a buy order and a sell order is easily derived from the solution for the sell order problem. It is the sum of the sell order value function and its reflection about zero:

$$V(e_t) = W(e_t) + W(-e_t)$$

[20] Note that this set changes across intervals because B_t is a continuous variable. Unfortunately, a constant grid would greatly complicate the quantitative analyses. Appendix B discusses the potential biases associated with this limitation.

III. INCENTIVE COMPATIBILITY

Since transaction cost minimization differs from sales price maximization, brokers may choose different order submission strategies than would their clients. The numeric solutions presented in Section IV below demonstrate these differences for several transaction cost measurement methods. Some transaction cost measurement methods do provide brokers with the proper incentives. This section presents some theoretical results on the incentive compatibilities of the various objective functions analyzed in this study.

Theorem 1 provides conditions under which uninformed brokers trading on behalf of uninformed liquidity traders will choose the proper strategies when transaction costs are measured relative to various price benchmarks. In general, brokers choose the optimal sales price maximization strategy when their decisions cannot predictably affect the measurement benchmark.[21] When trades are evaluated relative to a constant prespecified price, brokers choose the sales maximizing strategy because they cannot change the benchmark. When the benchmark is a post-interval price, brokers can change the realized benchmark value by deferring the execution of an order, but the expected benefits from deferral are zero because uninformed brokers cannot predict future prices. They therefore also choose the sales maximizing strategy.

A restatement of these results reveals that they may be somewhat surprising. Consider the risks that brokers face when using limit orders. Brokers whose trades are evaluated relative to a prespecified benchmark fear that the market will move away from their orders so that they will have to chase it. Such traders may reduce this risk by pricing their orders more aggressively, but this will decrease the trade price if the limit order does execute. Brokers whose trades are evaluated relative to a post-interval benchmark fear that their orders will execute when prices are moving through them. Such traders may reduce this risk by pricing their orders less aggressively, but this decreases the probability of filling the limit order. Theorem 1 shows that the optimal

---

[21] The result also assumes that the brokers are forced to internalize the costs of submitting orders and executing trades.

responses to these two seemingly different risks are the same: Both traders will use the same strategy (if the inter-interval mean μ is zero). The result can be understood intuitively by noting that the two risks are both related to the option value of the limit order, and by recalling that the traders analyzed in this study are all risk neutral.

Theorem 2 shows how the optimal solutions to the informed trading problem differ by the objective function specified. The sales price maximization objective and the value-added maximization objective produce identical optimal order submission strategies. The implementation shortfall minimization objective produces slightly different but generally similar optimal order submission strategies if the pricing error is expected to be small at the time the shortfall is computed for unfilled orders. When transaction costs are measured relative to a future price, traders choose strategies that are similar to those chosen by sales price maximizers only if the benchmark price is far in the future. If the benchmark price is the same price used to value an implementation shortfall, the problem has the same solution as the implementation shortfall problem. Finally, when transaction costs are measured relative to beginning-of-interval prices, traders will select less aggressive strategies than the optimal sales price maximizing strategy. The remainder of this section presents proofs and corollaries to these two theorems.

Theorem 1: Incentive Compatibility in the Liquidity Traders' Problem

In the liquidity traders' problem,

1. Sales price maximization and transaction cost minimization with respect to a prespecified price benchmark produce the same optimal order submission strategies.
2. Transaction cost minimization with respect to any post-interval price benchmark produces the same optimal order submission strategies as the sales price maximization when the inter-interval value innovation mean μ is zero.
3. Transaction cost minimization with respect to a beginning-of-interval price benchmark produces less aggressive optimal order submission strategies (when the inter-interval mean μ is zero) than does sales price maximization.

Proof

The proof proceeds by transforming the various valuation functions into comparable expressions. Before starting, recall from Section II.B that the uninformed trader believes price evolves according to the following process:

$$P_{t+1} = P_t + \varepsilon_t + \theta_t \tag{23}$$

where the intra-interval price innovation ε_t is distributed with zero mean and the inter-interval price innovation θ_t is independently distributed with mean μ so that

$$\mathrm{E}P_{t+s} = P_t + s\mu \,. \tag{24}$$

Constant Prespecified Benchmark: Proof of Point 1.

The value function for the liquidity trader's trading problem with a constant prespecified price benchmark P_o is given by

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_o - K_L - C_L & \text{limit order filled} \\ M_t - P_o - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{25}$$

and

$$W_T(P_T) = \mathrm{E}\left(M_T - P_o - K_M - C_M\right). \tag{26}$$

Add $P_o - \mathrm{E}P_{t+1}$ to both sides of (25) and (26), let

$$W_t^*(P_t) = W_t(P_t) - P_t + P_o \tag{27}$$

and take expectations to obtain

$$W_t^*(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}^*(P_{t+1}) - K_L + \mu & \text{unfilled limit order} \\ L_t - P_t - \varepsilon_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) + \mu & \text{no order submitted} \end{bmatrix} \tag{28}$$

and

$$W_T^*(P_T) = \mathrm{E}\left(M_T - P_T - K_M - C_M\right). \tag{29}$$

The term ε_t remains in Line 2 of (28) because its expectation, conditional on the execution of a limit order, is not zero. Since expression (28) does not depend on the prespecified price P_o, all problems with constant prespecified price benchmarks have the same solution. These problems include the sales price maximization problem for which $P_o = 0$ and the implementation shortfall problem for which P_o is generally specified as the midspread price at the time of order submission.

Post-Interval Benchmarks: Proof of Point 2.

Let $P_{t+s} + c$ be a post-interval price benchmark where $t+s$ indicates time at the end of the interval (t+) or later. The constant c may be set to $-\frac{S}{2}$, 0, or $\frac{S}{2}$ to respectively indicate a bid, quote midpoint, or ask price. The valuation function for this benchmark is given by

$$W_t(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_{t+s} - c - K_L - C_L & \text{limit order filled} \\ M_t - P_{t+s} - c - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{30}$$

and

$$W_T(P_t) = E(M_T - P_{T+s} - c - K_M - C_M). \tag{31}$$

Let

$$W_t^*(P_t) = W_t(P_t) + s\mu + c \tag{32}$$

and evaluate the expectations in (30) and (31) to obtain

$$W_t^*(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}^*(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t - \varepsilon_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{33}$$

and

$$W_T^*(P_T) = \text{E}\,(M_T - P_T - K_M - C_M). \tag{34}$$

When μ is zero, (33) and (34) are the same as (28) and (29) so that the two problems have the same solution.

Beginning-of-Interval Price Benchmark: Proof of Point 3.

Let $P_t + c$ be a beginning-of-interval price benchmark. The constant c may be set to $-\frac{S}{2}$, 0, or $\frac{S}{2}$ to respectively indicate a bid, quote midpoint, or ask price. The valuation function for this benchmark is given by

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t - c - K_L - C_L & \text{limit order filled} \\ M_t - P_t - c - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{35}$$

and

$$W_T(P_t) = E(M_T - P_T - c - K_M - C_M). \tag{36}$$

Let

$$W_t^*(P_t) = W_t(P_t) + c \tag{37}$$

to obtain

$$W_t^*(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}^*(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{38}$$

and

$$W_T^*(P_T) = \text{E}\left(M_T - P_T - K_M - C_M\right). \tag{39}$$

Transaction cost minimization with respect to a beginning-of-interval price benchmark produces different optimal order submission strategies than does sales price maximization because (38) and (28) differ: Expression (28) includes terms involving μ and ε_t. The conditional expectation (given a limit order execution) of ε_t is positive because sell limit orders execute, on average, when intraday prices rise. It is an increasing function of the limit price L_t because distant limit orders only execute when intraday prices rise substantially. Since ε_t enters Line 2 of (28) with a negative sign, it makes limit orders relatively less attractive than market orders, and it makes closely placed limit orders relatively more attractive than distant limit orders. Sales price maximizers will therefore choose more aggressive order strategies than will transaction cost minimizers who face a beginning-of-interval price benchmark.

Corollary 1.1

The optimal solution strategy for the uninformed liquidity traders' sales price maximization problem depends only on the time until the deadline and not also on the current price.

Proof

Recall from Section II.C that

$$EM_t = P_t - (1-\pi)\tfrac{1}{2}S \tag{40}$$

and let

$$\begin{aligned} L_t^* &= L_t - A_t \\ &= L_t - \left(P_t + \tfrac{1}{2}S\right) \end{aligned} \tag{41}$$

be the limit price expressed relative to the current ask so that (28) and (29) can be expressed as

$$W_t^*(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}^*(P_{t+1}) - K_L + \mu & \text{unfilled limit order} \\ L_t^* + \tfrac{1}{2}S - \varepsilon_t - K_L - C_L & \text{limit order filled} \\ -(1-\pi)\tfrac{1}{2}S - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) + \mu & \text{no order submitted} \end{bmatrix} \tag{42}$$

and

$$W_T^*(P_T) = -(1-\pi)\tfrac{1}{2}S - K_M - C_M\,. \tag{43}$$

Since the argument P_t in the valuation function W^* does not appear in lines 2 and 3 of (42) or on the right hand side of (43), and since the feasible set does not depend on P_t, the valuation function W^* does not depend on P_t. Only one state variable is needed to describe this transformed dynamic program. Since the operations necessary to obtain this transformed dynamic program are not dependent on the set of feasible strategies, the transformed dynamic program and the original dynamic program have the same optimal submission strategies.

Corollary 1.2

When μ is zero, the ex ante expected value of the valuation function given a post-interval benchmark is the same as its ex ante expected value given a prespecified price benchmark if the two benchmarks have the same relation to the midspread price.

Proof

When μ is zero, Theorem 1 shows that both problems have the same optimal solution because they both can be transformed to the same problem ((28) and (29) or (33) and (34)). The relation between the optimized valuation function for the untransformed prespecified price benchmark problem and the common transformed problem is given by (27). The analogous relation for the post-interval price benchmark problem is given by (32). Setting these expressions equal yields

$$W_t^o(P_t) - P_t + P_o = W_t^f(P_t) + s\mu + c \tag{44}$$

where $W_t^o(P_t)$ and $W_t^f(P_t)$ are respectively the valuation functions for the untransformed prespecified price benchmark problem and the untransformed post-interval (future) price benchmark problem. Let $P_o = P_r + c$ so that the prespecified price bears the same relationship to the midspread price as does the post-interval benchmark price, and let the subscript r indicate some interval before the start of the problem. Taking expectations as of interval r of both sides of (44) yields

$$\mathrm{E}W_t^o(P_t) - r\mu = \mathrm{E}W_t^f(P_t) + s\mu \tag{45}$$

which proves equality when μ is zero.

Theorem 2: Incentive Compatibility in the Informed Traders' Problem

In the informed traders' problem,

1. Sales price maximization and value-added maximization produce identical optimal order submission strategies.
2. Transaction cost minimization with respect to a post-interval price benchmark produces the same optimal order submission strategy as the implementation shortfall minimization method does when the price benchmark is the price used to price the shortfall.

3. Implementation shortfall minimization also produces the same optimal order submission strategy as sales price maximization when the expected pricing error $\hat{e}_{T+1}$ at the time the shortfall is priced, is very small.
4. Transaction cost minimization with respect to a post-interval price benchmark that is set at *s* intervals after the trade produces the same optimal order submission strategy as sales price maximization if *s* is large so that the pricing error $\hat{e}_{t+s}$ is very small.
5. Transaction cost minimization with respect to a post-interval price benchmark that is set at *s* intervals after the trade produces less aggressive order submission strategies than does sales price maximization if *s* is small.
6. Transaction cost minimization with respect to a beginning-of-interval price benchmark produces the least aggressive optimal order submission strategies.

Proof

The proof proceeds by transforming the various valuation functions into comparable expressions. Before starting, recall from Section II.B that the informed trader believes price evolves according to the following processes:

$$P_t = V_t + e_t \tag{46}$$

$$V_{t+1} = V_t + \varepsilon_t + \theta_t \tag{47}$$

where e_t is a pricing error, the intra-interval value innovation ε_t is distributed with zero mean and the inter-interval value innovation θ_t is distributed with mean μ. Recall also that the informed trader's estimate of the current pricing error is $\hat{e}_t$ and his expected estimate of e_{t+1} is given by

$$\hat{e}_{t+1} = \phi\hat{e}_t \ . \tag{48}$$

These assumptions imply that

$$P_{t+1} - P_t = e_{t+1} - e_t + \varepsilon_t + \theta_t \tag{49}$$

$$V_{t+s} = P_t - e_t + \sum_{i=1}^{s} \varepsilon_{t+i} + \sum_{i=1}^{s} \theta_{t+i} \tag{50}$$

and

$$P_{t+s} = P_t - e_t + \sum_{i=1}^{s} \varepsilon_{t+i} + \sum_{i=1}^{s} \theta_{t+i} + e_{t+s} \quad . \tag{51}$$

Implementation Shortfall Minimization

The valuation function for the implementation shortfall minimization problem (actually the maximization of the negative of the shortfall) is given by

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_o - K_L - C_L & \text{limit order filled} \\ M_t - P_o - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{52}$$

and

$$W_{T+1}(P_{T+1}) = P_{T+1} - P_o \tag{53}$$

where P_o is the prespecified price. Add $-P_t + P_o + \hat{e}_t$ to both sides of (52) and add $0 = -P_{t+1} + P_{t+1}$ to Rows 1 and 4 to get:

$$W_t(P_t) - P_t + P_o + \hat{e}_t = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - P_{t+1} + P_{t+1} - P_t + P_o + \hat{e}_t - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) - P_{t+1} + P_{t+1} - P_t + P_o & \text{no order submitted} \end{bmatrix} \tag{54}$$

Substitute (49) to get

$$= \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - P_{t+1} + \varepsilon_t + \theta_t + e_{t+1} - e_t + \hat{e}_t + P_o - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) - P_{t+1} + \varepsilon_t + \theta_t + e_{t+1} - e_t + \hat{e}_t + P_o & \text{no order submitted} \end{bmatrix} \tag{55}$$

Take expectations to get

$$= \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - P_{t+1} + \varepsilon_t + \mu + \hat{e}_{t+1} + P_o - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) - P_{t+1} + \mu + \hat{e}_{t+1} + P_o & \text{no order submitted} \end{bmatrix} \tag{56}$$

Let

$$W_t^*(P_t) = W_t(P_t) - P_t + P_o + \hat{e}_t \tag{57}$$

and subtract $E\varepsilon_t = 0$ from both sides to get the transformed valuation function

$$W_t^*(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}^*(P_{t+1}) + \mu - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - \varepsilon_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) + \mu & \text{no order submitted} \end{bmatrix} \tag{58}$$

The transformed terminal value condition is

$$W_{T+1}^*(P_{T+1}) = \hat{e}_{T+1} \ . \tag{59}$$

Sales Price Maximization

The transformed valuation function for the sales price maximization specification is almost identical to that of the implementation shortfall minimization specification. The dynamic part of the objective function is the same as (58) because a prespecified benchmark price of $P_o = 0$ implies sales maximization. The transformed terminal value condition differs, however, because the untransformed condition differs. The untransformed terminal value condition for the sales price maximization specification is

$$\begin{aligned} W_{T+1}(P_{T+1}) &= \text{E}V_{T+1} \\ &= P_{T+1} - \hat{e}_{T+1} \end{aligned} \tag{60}$$

The transformed terminal condition is

$$W_{T+1}^*(P_{T+1}) = 0 \ . \tag{61}$$

Value-Added Maximization

The value-added maximization valuation function is

$$W_t(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - \hat{V}_{T+1} - K_L - C_L & \text{limit order filled} \\ M_t - \hat{V}_{T+1} - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{62}$$

with terminal condition

$$W_{T+1}(P_{T+1}) = 0 \, . \tag{63}$$

Substituting (50) and taking expectations yields the transformed valuation function

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - \varepsilon_t - (T+1-t)\mu - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - (T+1-t)\mu - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{64}$$

Add $(T+1-t)\mu$ to both sides and let

$$W_t^*(P_t) = W_t(P_t) + (T+1-t)\mu \tag{65}$$

to get

$$W_t^*(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}^*(P_{t+1}) + \mu - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - \varepsilon_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) + \mu & \text{no order submitted} \end{bmatrix} \tag{66}$$

This transformation leaves the terminal condition unchanged.

Transaction Cost Minimization with Post-Interval Price Benchmark P_{T+1}

The valuation function for the post-interval price benchmark P_{T+1} is given by

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_{T+1} - K_L - C_L & \text{limit order filled} \\ M_t - P_{T+1} - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{67}$$

with terminal condition

$$W_{T+1}(P_{T+1}) = 0\,. \tag{68}$$

Substituting (51) and taking expectations yields the transformed valuation function

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - \varepsilon_t - (T+1-t)\mu - \hat{e}_{T+1} - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - (T+1-t)\mu - \hat{e}_{T+1} - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{69}$$

Add $(T+1-t)\mu + \hat{e}_{T+1}$ to both sides and let

$$W_t^*(P_t) = W_t(P_t) + (T+1-t)\mu + \hat{e}_{T+1} \tag{70}$$

to get

$$W_t^*(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}^*(P_{t+1}) + \mu - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - \varepsilon_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - K_M - C_M & \text{market order filled} \\ W_{t+1}^*(P_{t+1}) + \mu & \text{no order submitted} \end{bmatrix} \tag{71}$$

The transformed terminal condition is

$$W_{T+1}^*(P_{T+1}) = \hat{e}_{T+1} \ . \tag{72}$$

Transaction Cost Minimization with Post-Interval Price Benchmark P_{T+s}

The valuation function for the post-interval price benchmark is given by

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_{t+s} - K_L - C_L & \text{limit order filled} \\ M_t - P_{t+s} - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{73}$$

with terminal condition

$$W_{T+1}(P_{T+1}) = 0 \, . \tag{74}$$

Substituting (51) and taking expectations yields the transformed valuation function

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t + \hat{e}_t - \varepsilon_t - s\mu - \hat{e}_{t+s} - K_L - C_L & \text{limit order filled} \\ M_t - P_t + \hat{e}_t - s\mu - \hat{e}_{t+s} - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{75}$$

Transaction Cost Minimization with a Beginning-Of-Interval Price Benchmark

The valuation function for this case requires no transformation. It is

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L & \text{unfilled limit order} \\ L_t - P_t - K_L - C_L & \text{limit order filled} \\ M_t - P_t - K_M - C_M & \text{market order filled} \\ W_{t+1}(P_{t+1}) & \text{no order submitted} \end{bmatrix} \tag{76}$$

with terminal condition

$$W_{T+1}(P_{T+1}) = 0 \ . \tag{77}$$

Summary of Proof

Points 1 through 4 of the theorem are proven by inspecting the transformed valuation functions, which are summarized in Table 3. Transaction cost minimization with respect to a post-interval price benchmark set s intervals after the trade produces less aggressive order submission strategies than does sales price maximization when s is small (Point 5) due to the term $-\hat{e}_{t+s}$ that appears in lines 2 and 3 of the transformed valuation function for the former specification. This term, which is always negative (for an informed sale) makes order execution less attractive. Transaction cost minimization with respect to a beginning-of-interval price benchmark produces the least aggressive optimal order submission strategies (Point 6) because the value function does not include the positive term $\hat{e}_t$ in lines 2 and 3 and because it does not include the term $-\varepsilon_t$ in Line 2. The omission of the former term makes all order executions less attractive. The omission of the latter term makes limit orders placed far from the market more attractive than limit orders placed close to the market because its expected value (given a limit order execution) is negative and its absolute size decreases the further from the market the limit order is placed.

IV. NUMERIC SOLUTIONS AND COMPARATIVE STATIC RESULTS

The numeric solutions provided in this section are interesting both for practical and theoretical reasons. For practitioners who accept that the model assumptions and the chosen parameter values reasonably represent their problems, the quantitative results provide specific solutions that may help them form their order submission strategies. Theoretical interest in these results comes from the difficulty of obtaining formal expressions for the model comparative statics. Although most effects from varying the parameters can be predicted by examining the model specification, they cannot be expressed formally because the various integrals that appear in the objective function and the recursive nature of the dynamic problem complicate the analysis. The numeric solutions allow us to explore the comparative statics of the models.

This section starts with a discussion of how the models are solved. It then examines solutions for each of the three stylized problems. For each problem, a baseline set of parameters is specified and the corresponding solution is obtained. Model comparative statics are examined by varying parameters in the baseline specification and comparing the resulting solutions to the baseline solution.

A. SOLUTION METHODS

The various valuation functions specified in this study for the prespecified uninformed liquidity trader problem and for the informed trader problem all can be expressed in the following form:

$$W_t(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}(P_{t+1}) + C_1 & \text{unfilled limit order} \\ L_t - P_t - \varepsilon_t + C_2 & \text{limit order filled} \\ M_t - P_t + C_3 & \text{market order filled} \\ W_{t+1}(State_{t+1}) + C_4 & \text{no order submitted} \end{bmatrix} \quad (78)$$

where the various C_j are constants that may depend on t but not on the price state variable or on variables that generate prices. In some specifications, the price state variable and/or the value-innovation ε_t do not appear.

To simply the calculations and the presentation of the results, each problem is restated to make the feasible limit order space and the state space stationary. Using (11), (78) is transformed to

$$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) + C_1 & \text{unfilled limit order} \\ L_t^* - \varepsilon_t + C_2^* & \text{limit order filled} \\ -(1-\pi)\tfrac{1}{2}S + C_3 & \text{market order filled} \\ W_{t+1}(P_{t+1}) + C_4 & \text{no order submitted} \end{bmatrix} \qquad (79)$$

where $L_t^* = L_t - A_t$ is the limit price expressed relative to the current ask price and $C_2^* = C_2 + \tfrac{1}{2}S$. The maximization is then done over the stationary variable L_t^* instead of the non-stationary variable L_t. Likewise, the state variable for the informed trader's problem with the reservation price constraint, is transformed from P_t to $PMR_t = P_t - P^R$ to make the state variable stationary. (The state variable in the value-motivated traders' problem, $\hat{e}_t$, is already stationary.)

Several numeric integrals must be computed to evaluate these valuation functions. In each problem, the probability that a limit order executes must be computed for all feasible limit order strategies. This probability and its complement are used to compute the expected values of the constants in lines 1 and 2 of (79). In addition, for those problems that require it, the partial expectation of the value innovation ε_t (line 2) must be computed for all outcomes that lead to a limit order execution. Finally, for those problems in which the valuation function depends on price, two additional integrals also must be computed. The first is the partial expectation for the valuation function for limit order submissions that fail to execute (Line 1). The other integral is the unconditional expectation of the valuation function (Line 4). Appendix A describes the methods used to compute these numeric integrals.

Those problems that have terminal conditions (the uniformed liquidity traders' problem and the informed traders' problem) are solved by backwards induction starting from their terminal conditions. At each interval, the valuation function is evaluated for each feasible strategy given the valuation function at the next interval. The best strategy is then selected, and its associated value is the value of the problem in that interval.

The infinite period value-motivated traders' problem is solved by finding a strategy for each state variable level (the current pricing error estimate) that gives a fixed point in the function space for the valuation function $W(\hat{e}_t)$. The problem is easier than it seems since (as noted above in Section II.D) the optimal strategy depends only on $\hat{e}_t$ and not also on $W(\hat{e}_t)$. The optimized valuation function is found by assuming a discrete lattice to represent the continuous variable $\hat{e}_t$. The maximized valuation function (21) can then be expressed as a linear matrix equation that simultaneously describes that valuation function and the numeric integral necessary to evaluate the unconditional expectation of the valuation function. The fixed point solution is found by solving this equation.[22]

B. THE LIQUIDITY TRADER PROBLEM

The first set of results (Table 4) explores how spreads and volatility affect optimal order submission strategies. Since these variables are the most important determinants of order submission strategy, solutions are presented for a variety of scenarios that practitioners may face.

Baseline Parameters

The baseline solution assumes that the liquidity trader maximizes the difference between the net sales price and the prevailing bid at the time the decision to trade was made. Theorem 1 of Section III shows that this objective produces the same optimal order submission strategies as the sales price maximization objective. This objective was chosen for the baseline solution because the optimized valuation function lends itself to two easy interpretations: It is the expected portfolio implementation shortfall, and, when μ is zero, it is equal to the expected price improvement measured relative to the end-of-interval bid (Corollary 1.2).

The execution and submission costs are set to values that practitioners likely face for 1,000 share orders. The execution costs per share are set to 3 cents for both order types. Submission costs are set to zero for market orders and to 0.2 cents per share for limit orders. The

---

[22] This equation is

$$\mathbf{w} = \mathbf{A}\mathbf{w} + \mathbf{c}$$

where **w** is the vector of value function values corresponding to the lattice for $\hat{e}$, **A** is a matrix of integration weights, probability densities and discount factors, and **c** is a vector of the maximized values of the right most term of (21).

limit order submission cost is primarily due to the surveillance required to manage limit orders that do not fill. At 0.2 cents per share, this cost is 2 dollars per interval.

The market order price improvement parameter is set to 50% of the half-spread. This fraction is suggested by empirical results in Harris and Hasbrouck (1996). Although they show that the price improvement rate varies somewhat with the bid/ask spread, the parameter is held constant in Table 4 so that the effects of changing the bid/ask spread will be unambiguous.

The two degrees of freedom parameters that appear in the beta distribution assumed for the limit order fill probability distribution are both set to 1. The resulting distribution is uniform.[23]

The degrees of freedom parameter that appears in the Inverted Gamma distribution for volatility (and hence in the implied Student-*t* distribution for unconditional price changes) is set to 3. This low value ensures that the probability of large price changes will be large. It is on the lower side of empirical results reported by Praetz (1972) for daily price changes, and on the higher side of unpublished empirical results obtained by the author for higher frequency data.

The inter-interval drift is set to zero to reflect its approximate value over short intraday intervals. This assumption is consistent with empirical results found in Harris (1986b) and elsewhere. These results show that the average intraday return on stocks has been near zero. Most of the positive average holding returns on stocks have accrued while the markets are closed.

The discrete tick used to specify the feasible limit orders is set to 1/8 dollar. This value is used for most stocks traded at US and Canadian stock markets.

Quantitative Results

Table 4 presents solutions to the uninformed liquidity trader problem for various combinations of price change volatility and spread. All other parameters are held constant.

---

[23] These parameters can be estimated using data on limit order execution frequencies. This exercise is left for future research. Although the results will probably vary with stock volatility, firm size and bid-ask spread, the uniform distribution is used in this analysis so that the effects of changing bid-ask spreads and changing volatilities will be unambiguous.

Solutions are computed for five values of the intra-interval price change standard deviation. These values represent a wider variety of problems than might be immediately obvious. Since the price change standard deviation is the product of the return standard deviation times the price level, the solution for a 10 dollar security with a 2 percent daily return standard deviation with decisions made over daily intervals is identical to the solution for a 40 dollar security with a 0.5 percent half-hourly return standard deviation with decisions made over half-hour intervals. Both problems have a price change standard deviation over their respective intervals of 20 cents.

The numeric solutions show that orders should be placed close to the market when the deadline is near (Table 4, Panel A). When the deadline is distant, orders should be placed farther from the market. Though these orders will be filled less often (Panel B), the trader will usually be able to avoid using a market order (Panel C). The dynamic limit order strategy produces higher sales price variance, however (Panel D). The value of the trading program is greatest for distant deadlines (Panel E), and most of the value can be obtained within just a few intervals (Panel F).

The value of the trading program generally is lower for high volatilities (Panel E, Rows 3-10). High volatilities make limit order strategies less attractive because volatility increases limit order option values. Low volatilities allow the trader to submit orders close to the market without giving up too much option value. Such orders often fill before the market can run too far away in the other direction. However, if little time remains until the deadline, the value of the program is smaller for low volatilities that are small relative to the discrete tick (Panel E, Row 1). In such cases, the optimum limit order price for a given volatility is not available on the discrete lattice and the volatility is not great enough to ensure a favorable execution at the nearest tick before the deadline.

High spreads increase the benefits of using limit orders. Given the assumed uniform fill probability distribution, an increase in the spread increases the probability that any given limit order will fill. Since limit orders in this model execute only when the limit price is between the bid and offer (inclusive), the expected value of the program is never greater than the spread (Panel E).

For a given volatility, the effect of the spread on the order submission strategy depends on the time remaining until the trade deadline. If little time remains, the limit price is set close to the prevailing ask when spreads are wide: Traders want to ensure that their limit orders fill so that they benefit from the wide spread instead of paying it should they be forced to use a market order. If the deadline is distant, the strategy is less aggressive when the spread is wide: The cost of failing to trade is lower because a more aggressively priced order placed later will probably still fill at a good price. When the aggressiveness of the strategy is measured relative to the spread midpoint, less aggressive strategies are uniformly chosen when the spread is wider: Wider spreads increase the probability that any given limit order will fill.

Table 5 illustrates the effects of changing other parameters. The first column presents a baseline solution (the same as in Column 9 of Table 4) to which other solutions are compared. The baseline parameter set includes all those mentioned above. In addition, the bid/ask spread is set at 25 cents and the interval value innovation standard deviation is set at 40 cents. The other columns present solutions for which the value of one parameter is changed.

Traders are substantially less aggressive when their trading is evaluated relative to beginning-of-interval prices than prespecified (or equivalently end-of-interval) prices (Column 2). They place their orders far from the market and wait for an extreme price change that will make them look good when they execute. The value of this trading problem is much higher than the value of the baseline problem, but the difference in value is not real because it does not accrue to the portfolio. It arises because the benchmark is reset each interval: the trader is not penalized for executions that result after chasing a fall in prices.[24]

The results in Columns 3-5 illustrate the easily predictable effects of the execution difficulty parameters. A decrease in expected market order price improvement makes market orders less attractive so that limit order strategies become more aggressive (Column 3). A shift of mass in the execution fill probability density towards higher prices makes limit sell orders

---

[24] An increase in volatility will make these traders less aggressive and further increase the value of their problem.

easier to execute and thereby favors limit order strategies (Column 4).[25] Limit order strategies become more aggressive because the higher probability of quick execution decreases the option value of the order. Finally, limit orders are not used when they can only be executed at dealer bid prices (Column 5). When the trading mechanism requires the seller to trade at the bid, limit orders cannot benefit from the spread. In such markets, sellers use market orders to avoid losses due to limit order option values.

The order submission and execution cost parameters also have easily predictable effects. If both limit and market order submission costs are high, the limit order strategy should be aggressive. It should also be aggressive if the limit order submission cost is greater than the market order submission cost. If execution costs are equal for both order types, these costs should not affect order submission strategy because the liquidity trader must trade and incur them in any event. If the commission for limit orders is higher than for market orders, as is sometimes the case, limit order submission strategies should be less aggressive.

Columns 6 and 7 illustrate two of these effects: Market orders are more attractive when their execution cost falls to 2 cents per share while keeping the limit order cost at 3 cents (Column 6). In this baseline problem, this small change—which might be due to payment for order flow—only affects the value of the problem and not the order submission strategy. If the tick were smaller, the effect on strategy would be apparent. Second, traders submit limit orders at more aggressive prices when their submission cost increases from 0.2 cents per share to 0.4 cents per share (Column 7).

The value of the dynamic trading program increases when the tick size decreases to one cent (Column 8, Panel C). The smaller tick size allows more precise strategies to be used (Panel A). The increase in value, however, is remarkably small. The coarse discrete feasible set in the baseline problem does not significantly constrain the solution.[26]

---

[25] The mass is shifted by adjusting the two beta distribution degrees of freedom parameters so that the mean shifts while holding the variance constant.

[26] Although this result suggests that a decrease in tick size, such as that proposed in the SEC Division of Market Regulation *Market 2000 Study* (1994), would not benefit limit traders, it must be noted tick size has other effects on

Trading strategies become less aggressive when the inter-interval mean price change increases to 0.5 cents per share (Column 9). When the interval is a day, 0.5 cents per share per day represents a compounded expected return (presumably in excess of some alternative) of 6.5 percent per year (252 trading days) for a 20 dollar security. The positive drift makes trading less aggressive because it benefits the trader who trades near the deadline.

Finally, for a given level of unconditional volatility, the value of the program increases when value innovations are normally distributed (Column 10). The stochastic volatility assumed in the baseline solution increases the option value of limit orders by increasing the probability of extreme events, and thereby decreases the value of the program.

Comments

The values of all program specifications described in this study increase with the number of intervals to the trading deadline. Although this result is not surprising—relaxing a constraint never harms the solution to an optimization problem—it is useful to interpret the result.

A distant trading deadline benefits the trader two ways. First, when a trader places a limit sell order at the ask or below, for a given midspread price, time increases the probability that a counter-party will arrive who wants to trade. Formally in this model, the probability of trading increases the more often the trader is allowed to draw from the fill probability distribution. Second, time allows the limit order trader to pursue an order strategy that minimizes the option value of his order. Traders can lower limit order option values by placing them far from the market. Although such placements lower the probability of executing in any given interval, a distant trading deadline increases the probability that the trade will eventually take place by increasing the number of opportunities for an extreme price change.

The price uncertainty that results from using limit order strategies may cause risk averse traders to favor more aggressive strategies that will fill their orders quicker. Since the benefits of a long deadline are declining, the cost of using more aggressive strategies may be small relative

---

traders than those modeled here. For example, empirical results in Harris (1994) suggest that bid-ask spreads would decrease if tick size were decreased. Holding other factors constant, a decrease in spreads would harm limit order traders.

to the associated reduction in risk. A risk tolerant trader therefore can be encouraged to pursue more aggressive strategies by artificially restricting the trading deadline.[27]

C. The Informed Trader Problem

The informed trader problem introduces three new features that are not present in the liquidity trader's problem. The first feature is an information advantage that decays with time. The second feature is the option not to trade. The final feature is the reservation price constraint.

This section examines the effects of these features on the optimal informed trading strategy. Since the parameters that are also present in the liquidity trader's problem have similar effects in this problem, they are not analyzed further.

The analysis first considers properties of the solution that are related to the trader's private information. To keep the analysis simple and to facilitate comparisons to the liquidity trader problem results, the reservation price constraint is not initially imposed. The resulting program has only one state variable—time. The effect of the reservation price constraint on the solution is examined at the end of this section.

Column 1 of Table 6 presents a baseline solution to which the other solutions will be compared. The baseline specification assumes the informed trader maximizes the expected value-added from trading. (Theorem 2 of Section shows that this objective is the same as maximizing the sales price.) The expected pricing error persistence parameter ϕ is assumed to be 70 percent per interval and prices are assumed to be 400 cents above estimated value at 10 intervals before the trading deadline. The values of all other parameters that appear in the problem are the same as in the baseline specification for the liquidity trader problem.

The baseline solution (Panel A) is substantially different from the baseline solution to the liquidity trader's problem. The strategy is most aggressive when the order is first submitted since the trader tries to quickly capitalize on his information. If the order does not execute, some

---

[27] Risk aversion can be indirectly incorporated into the model by increasing the submission cost for limit orders to reflect the greater price uncertainty they create. The incremental cost should vary by the limit order price since orders placed far from the market create more execution price uncertainty than do orders placed near the market.

information advantage is lost and the trader becomes less willing to incur transaction costs to profit from his advantage. By the end of ten intervals, the expected pricing error is only $400 \cdot 0.7^{10} = 11.3$ cents (Panel D).

Two dimensions of the trader's informational advantage—the size of the expected price fall and the expected rate at which it will occur—affect order strategy. When either variable is large, the order submission strategy will be aggressive. Columns 2 and 3 of Table 6 illustrate these effects. The trading strategy is more aggressive for a larger initial pricing error of 1,000 cents (Column 2) and for a smaller persistence rate of 0.5 (Column 3). In the latter specification, the expected pricing error eventually becomes so small (Panel D) that the trader stops trying to fill the order three intervals before the deadline.

When the pricing error persistence rate is high, the expected pricing error at the deadline remains economically significant. The trader's strategy near the deadline is then like the liquidity trader's strategy because he does not want the order to go unexecuted. When the persistence rate is 90 percent (Column 4), the limit order pricing strategy is no longer monotonic. At the beginning, the order submission strategy aggressively tries to capture the full value of the information advantage. As time passes, the strategy becomes less aggressive because the expected loss from failing to execute is decreasing. As the deadline approaches, the strategy becomes aggressive to ensure that the remaining expected advantage of 139 cents is not lost.

The end-of-interval and beginning-of-interval objective function specification of the informed trader problem produce less aggressive strategies than does the value-added maximization specification (Columns 5 and 6). These objective functions do not properly weight the depreciation of the trader's information advantage. The end-of-interval benchmark trader does not even bother to trade because the expected value of his benchmark is less than zero. If the probability fill distribution were more favorable or if the spread were wider, this trader would submit orders that would be less aggressive than in the baseline problem, but more aggressive than for the beginning-of-interval benchmark (results not shown).

The Effect of a Reservation Price Constraint

The reservation price constraint adds an additional state variable—current price slack—to the problem. The current price slack is the difference between the quotation midpoint and the reservation price at the beginning of the interval. When the slack is negative, orders must be set far from the market so that their prices are above the reservation price. Limit prices therefore must be greater than half the spread minus the current price slack. When the slack is less than half the spread, market orders are not feasible because they may execute at the bid price.

The probability that a price constraint will be binding in the future depends on the volatility in the price process. The optimal order submission strategy therefore depends on the volatility of price changes that occur between intervals, which has not been previously specified.[28] For the baseline specification, assume that the inter-interval price change standard deviation is 1 cent. This small number is appropriate for contiguous intervals that occur sequentially within the same trading session.[29]

The optimal solution to the baseline problem with a reservation price constraint appears in Table 7. Each block in the table presents the solution for a different number of intervals to the deadline and for various values of the current price slack. For large current price slack values the constraint is not binding so that the solution is identical to the unconstrained baseline problem (Table 6, Column 1). For lower values, the constraint is either binding or may likely be binding in the future. Since the unconstrained optimum order may no longer be feasible, the value of the constrained problem is lower. Orders are generally less aggressive for lower slack values, but the optimum order strategy is not always monotonic in the slack. It occasionally becomes more aggressive slightly above zero slack values. This inflection in the optimal strategy is an attempt to lock in profit before the expected price decrease causes the reservation price constraint to become binding. The inflection appears in this table at intervals 5, 6, 7, and 8 before the deadline. (The inflection is also present at interval 9, but for slack values that do not appear in the table.)

---

[28] In the previous problems examined in this section, the inter-interval price variance did not affect the solution because the traders are risk neutral and because the problems have no other asymmetries related to price.

[29] The inter-interval volatility is not set to zero for contiguous intervals because the results in Appendix A show that uncertainty in the pricing error estimates has the same effect on the problem as does inter-interval volatility.

To determine the effect of inter-interval volatility on the results, the inter-interval price change standard deviation was set to 40 cents and a new solution was obtained. The optimized value function declines (results not reported) as expected. In all other respects the solution is virtually identical to that reported in Table 7.

D. THE VALUE-MOTIVATED TRADING PROBLEM

The value-motivated trading problem has three new parameters that must be assigned values to complete the baseline specification. First, assume that the Brownian motion process that generates innovations to the pricing error process has a standard deviation of 20 cents per interval. Next assume that the pricing error innovations are uncorrelated with the value innovation process. These two parameters imply that the value-motivated trader believes price changes measured across intervals have a serial correlation of -0.034 when the pricing error persistence rate is 70 percent.[30] Finally, assume that the intertemporal discount rate is 0.05 percent per interval. If the interval is one day, this discount rate corresponds to an annual (252 trading days) compounded rate of 13.4 percent. (This last assumption only affects the value of the problem and not the optimal strategy.) The values of all other parameters that appear in the problem are the same as in the informed trader baseline specification.

Column 1 of Table 8 presents the baseline solution to which other solutions will be compared. The solution consists of a strategy (Panel A) and an associated optimized valuation

---

[30] The formula for the serial correlation is derived from the value-motivated traders' price evolution model. Their model implies

$$\begin{aligned}\Delta P_t &= \Delta V_t + \Delta e_t \\ &= \varepsilon_t + (\phi - 1)e_t + n_t\end{aligned}$$

so that

$$SCorr(\Delta P_t) = \frac{\mathrm{SCov}(\Delta P_t)}{\mathrm{Var}(\Delta P_t)} = \frac{(\phi-1)^2\phi\sigma_e^2 + (\phi-1)\sigma_N^2 + (\phi-1)\sigma_{\varepsilon N}}{\sigma_\varepsilon^2 + (\phi-1)\sigma_e^2 + \sigma_N^2 + 2\sigma_{\varepsilon N}}$$

where

$$\sigma_e^2 = \frac{\sigma_N^2}{1-\phi^2}.$$

function value (Panel B) for each level of the current expected pricing error. To obtain the numeric solutions, only a discrete lattice of expected pricing error values were considered. Appendix A describes the solution method.

The results show that value-motivated traders trade aggressively when the expected pricing error is large. When the expected pricing error is small, traders are less aggressive. When the expected pricing error is small and negative, value-motivated traders place sell orders far from the market to ensure profitable trades should the pricing error reverse. If the pricing error is very negative, they do not submit any orders because the probability that price will revert far enough to allow profitable trades is too small. The expected profits from submitting such distant orders are less than the costs of submitting them.

Value-motivated traders who simultaneously solve this trading problem for both sides of the market will sometimes submit orders on both sides of the market. When the pricing error is small and negative, they submit aggressive buy orders and less aggressive sell orders. When the pricing error is small and positive, they submit aggressively priced sell orders and less aggressively priced buy orders. The spread between their orders can be computed from the numeric results. When the pricing error is 0, value-motivated traders place sell orders 75 cents above the ask and buy orders 75 cents below the bid. Since the baseline solution bid/ask spread is 25 cents, the spread between these two orders is 175 cents. Treynor (1987) calls this spread the outside spread. It widens to 212.5 cents when the absolute pricing error is 20 cents, and widens further for greater pricing errors.

For a given pricing error, the value-motivated trader is more aggressive than the informed trader because the former wants to complete the trade quickly so that another trade can be done. For example, when the expected pricing error is about 60 cents with five intervals until the deadline, the informed trader places a limit order at the ask (Table 7, Panels A, Column 1, line 5). For the same expected pricing error, the value-motivated trader places a limit order at the spread midpoint (Table 8, Panel A, Column 1, at 60).

The value of the program is high for all values of the expected pricing error (Panel C). The high values reflect the cumulated present value of trading indefinitely. For very high values of the state variable, the program value increases approximately in linear relation to the state variable (results not shown). Very large pricing errors will rarely arise from the assumed pricing error process. The program values for these extreme pricing errors reflect the benefits of receiving a one-time windfall of very valuable information.

An increase in the standard deviation of the pricing error innovation process to 40 cents per interval increases the value of the program (Panel C, Column 2). Value increases because the average absolute pricing error increases. The trading strategy becomes less aggressive as the trader lies in wait for a larger pricing error (Panel A). The trading strategy also is less aggressive because the increase in volatility increases the option value of limit orders.

A decrease in the pricing error persistence rate to 0.5 decreases the value of the program (Panel C, Column 3). The smaller persistence rate decreases the expected benefits from trading in the next interval and it also decreases the unconditional volatility of the pricing error process. Large errors decay faster and fewer large errors arise so that less profit will be made. Surprisingly, the trader's optimal strategy does not depend on the persistence (Panel A). The invariance is due to the repeated nature of the problem and to the assumption that no reversion in the pricing error process occurs within the interval. (The pricing error is assumed to decay only between intervals.) Since the trader knows that he will return in the next interval, whether he trades or not, he is only concerned about profiting from the current pricing error. Formally, the invariance arises because the persistence parameter appears only in the value function and, as noted above, the optimal strategy does not depend on the value function.

An increase in the value innovation standard deviation to 80 cents per interval decreases the value of the program (Panel C, Column 4). This change increases the option value of limit orders but does not affect the pricing error process that is the source of the trading profits. The trader moves his orders further from the market to limit their option values (Panel A). They execute less frequently (Panel B) so that the value of the problem decreases.

A positive correlation of 0.8 between the value and pricing error innovations increases the value of the problem and makes the optimal strategy less aggressive (Column 5). The positive correlation increases the probability that limit orders will execute when there is an innovation in the pricing error process. Likewise, a decrease in the correlation to -0.8 decreases the value of the problem and makes the optimal strategy more aggressive (Column 6).

An increase in the intertemporal discount rate to 0.1 percent per interval (28.6 percent per year for daily intervals) decreases the valuation function but has no effect on the optimal strategy (Column 7). The decrease in the valuation function is due to the greater compounding rate. The strategy does not depend on the discount factor β because β does not appear in the expressions over which the optimization is conducted.

V. CONCLUSION

A. SUMMARY

This study formulates and solves some dynamic programs that represent common trading problems faced by investors. Three stylized problems are considered. The liquidity trader problem considers how an uninformed trader should trade. The informed trader problem considers how an informed trader should trade when he only receives a single signal about value. The value-motivated trader problem considers how a trader who continuously estimates security values should trade.

The results are highly intuitive. Traders are most aggressive when volatility is high and when their information advantages, if any, are large and decay quickly. Traders are patient when their deadlines are not pressing and when bid/ask spreads are wide. Most traders issue limit orders only if they can capture a portion of the bid/ask spread.[31] Limit orders should not be used when public limit orders can only be executed against dealer quotes: Traders cannot capture any of the spread in such markets. Value-motivated traders may use limit orders to preposition themselves to take advantage of any mean reversion that they may believe to be present in prices.

The numeric results suggest that most traders should place limit orders close to the market when they trade. Although it may sometimes be optimal for risk neutral traders to place orders far from the market (when deadlines are distant or when private information will not be revealed soon), the expected additional benefits from this strategy are very small. If monitoring

---

[31] A trader "captures" a part of the spread when using a limit order if the trade price is better than would have been obtained had the trader used a market order.

of open orders is expensive or if the trader is risk averse, distant order placement strategies will not be optimal. The only exception to this rule is for traders who believe that prices are mean-reverting. They may place limit orders far from the market to benefit if prices move far from fundamental values.

These results are generally consistent with empirical results presented in Harris and Hasbrouck's (1996) study of actual SuperDOT order transaction costs. For most stocks, the optimal order submission strategy for a single order that will not be replaced is a limit order placed at the quote or slightly inside it. Although this study of dynamic submission strategies does not impose a single order submission constraint, the optimal dynamic strategy often places its first order close to the market so that it has a high probability of executing.

These results also show that brokers will adopt portfolio value maximizing trading strategies when their transaction costs are measured relative to prespecified price benchmarks. The portfolio implementation shortfall transaction cost measurement method, popularized by Perold (1988), uses such prices. End-of-interval price benchmarks also work well when orders are not based on private information. Beginning-of-interval price benchmarks create distorted incentives for traders. Traders facing these benchmarks will trade less aggressively than they should. The volume-weighted average price benchmark is not analyzed in this study, but its properties are easily characterized using these results. This benchmark uses a price that, in a temporal sense, is between the beginning-of-interval and end-of-interval prices. It will therefore share some of the properties of both benchmarks. It differs from both, however, since traders can estimate its value before they set their orders. The volume-weighted average price benchmark therefore gives traders a valuable timing option. Their average estimated transaction costs will be lower if they can costlessly defer the execution of an order into an interval during which a new volume-weighted average price will be computed.

B. LIMITATIONS

The results obtained in this study apply primarily to small trades since the models assume that orders have no impact on prices. If orders do affect prices, traders should consider whether, when and how their orders should be broken up to obtain better executions. Although these questions are very important, they are beyond the scope of the present analysis, which is designed to investigate the relation between immediacy and the cost of liquidity. Since size clearly affects this relation, it should be the subject of further research, such as is presented in Bertsimas and Lo (1996). Despite this reservation, the results should still be interesting to practitioners since most institutional orders and almost all individual orders are small relative to daily volume. In addition, the intuition that these models formalize applies to trades of all sizes.

This analysis also does not adequately model price discreteness for some trading problems. The analysis imposes discreteness only on the set of feasible limit order prices and not also on the price process itself. As a consequence, results may be biased for securities whose price volatility over the interval is small relative to the tick size. The biases probably are not large for securities priced over 10 dollars that are monitored over daily intervals or for securities priced over 40 dollars that are monitored over 30 minute intervals.[32] A detailed discussion of the discreteness problem appears in Appendix B.

Finally, the models presented in this study treat the bid/ask spread and various parameters that determine execution difficulty as exogenous variables. These variables are actually endogenous variables since they depend on trader and dealer order submission strategies. For example, spreads may tighten if more traders choose to submit limit orders. A tightening of the spreads, however, makes limit order strategies less attractive. With additional effort, these models can be made into equilibrium models in which these endogenous variables are determined by other, presumably more exogenous variables like the flow of liquidity traders.[33]

---

[32] Although it is impossible to determine the magnitudes of discreteness-induced biases without a model of discrete prices, it is probably safe to assume that they are not large when the price change standard deviation of the interval is greater than one and one-half ticks. Expressed formally, this condition is

$$P\sigma\sqrt{D} > 1.5d$$

where P is price level, σ is the return standard deviation, D is the length of the interval, and d is the discrete tick. For $\sigma = 0.02$, $P = 40$, and $d = 1/8$, D must be greater than 5½ percent of a day, or about 30 minutes if the trading day is 10 hours long. For $P = 10$, D must be greater than 90 percent of a day.

[33] Cohen, Maier, Schwartz and Whitcomb (1981) provide a related example of a two-period equilibrium model for spreads.

This analysis was not done, however, because the qualitative results would be obvious from the comparative statics of our models, and because the quantitative results would depend completely on the assumptions made about exogenous variables like the number of uninformed traders. Although this study does not provide a complete equilibrium analysis, it does provide normative results about the demand and supply of immediacy that furthers our understanding of that equilibrium.

REFERENCES

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

Beebower, G. 1989. "Evaluating Transaction Cost." Pp. 137-50 in *The Complete Guide to Securities Transactions*, ed. W. Wagner. New York: John Wiley & Sons.

Berkowitz, S. A., D. E. Logue, and E. A. Noser, Jr. 1988. "The Total Cost of Transactions on the NYSE." *Journal of Finance* 41, 97-112.

Bertsimas, D. and A. Lo. 1996. "Optimal Control of Execution Costs," MIT working paper.

Biais, B., P. Hillion and C. Spatt. 1995. "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse." *Journal of Finance* 50, 1655-89.

Bronfman, C. 1992. "Limit Order Placement, the Supply of Liquidity, and Execution Risk: An Empirical Investigation." Unpublished Paper, University of Arizona.

Chakravarty, S. and C. W. Holden. 1995. "An Integrated Model of Market and Limit Orders." *Journal of Financial Intermediation* 4, 213-41.

Cohen, K. J., D. Maier, R. Schwartz and D. Whitcomb. 1981. "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread." *Journal of Political Economy* 89, 287-305.

Copeland, T. E. and D. Galai. 1983. "Information Effects on the Bid-Ask Spread." *Journal of Finance*, 38, 1457-69.

Handa, P. and R. A. Schwartz. 1993. "Limit Order Trading." Unpublished Paper, Stern School of Business.

Harris, L. 1986a. "Cross-Security Tests of the Mixture of Distribution Hypothesis." *Journal of Financial and Quantitative Analysis* 21, 39-46.

Harris, L. 1986b. "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns." *Journal of Financial Economics* 16, 99-117.

Harris, L. 1987. "Transaction Data Tests of the Mixture of Distributions Hypothesis." *Journal of Financial and Quantitative Analysis* 22, 127-41.

Harris, L. 1990a. "Liquidity, Trading Rules, and Electronic Trading Systems." *New York University Salomon Center Monograph Series in Finance*, Monograph 1990-4.

Harris, L. 1990b. "Estimation of Stock Price Variances and Serial Covariances from Discrete Observations." *Journal of Financial and Quantitative Analysis* 25, 291-306.

Harris, L. 1994. "Minimum Price Variations, Discrete Bid/Ask Spreads and Quotation Sizes." *Review of Financial Studies* 7, 149-78.

Harris, L. and J. Hasbrouck. 1996. "Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy." *Journal of Financial and Quantitative Analysis* 31, 213-231.

Kumar, P. and D. Seppi. 1993. "Limit and Market Orders with Optimizing Traders." Unpublished Paper, Carnegie Mellon University.

Perold, A. 1988. "The Implementation Shortfall: Paper Versus Reality." *Journal of Portfolio Management* 14, 4-9.

Praetz, P. D. 1972. "The Distribution of Price Changes." *The Journal of Business* 45, 49-55.

Shepp, L. A. 1979. "The Joint Density of the Maximum and its Location for a Wiener Process with Drift." *Journal of Applied Probability* 16, 423-27.

Treynor, J. 1987. "The Economics of the Dealer Function." *Financial Analysts Journal*, 27-34.

US Securities and Exchange Commission, Division of Market Regulation, 1994. *Market 2000: An Examination of Current Equity Market Developments.* Washington, DC: US Government Printing Office.

APPENDIX A

MATHEMATICAL NOTES ON THE NUMERIC INTEGRATIONS

This appendix explains how to evaluate the four numeric integrals that appear in the objective functions. The integrals are mostly three-fold convolutions (mixtures) of probability distributions. The first distribution is a conditional probability statement about a first passage event for a continuous Brownian motion with constant volatility. This analytic expression is then numerically integrated over the fill probability distribution that specifies the probability that an order will execute for a given passage event. The resulting numeric integral is then numerically integrated over the Inverted Gamma distribution for interval volatilities.

Section A derives passage event properties for symmetric zero drift diffusion processes. These properties are then used in Section B to compute expressions that involve the limit order fill probability distribution. Finally, Section C describes how to use these expressions to compute the numeric integrals that appear in the objective functions of the dynamic trading problems.

A. PASSAGE PROPERTIES OF SYMMETRIC DIFFUSION PROCESSES

This subsection derives probabilities, partial probability density functions and partial expectations associated with passage events for symmetric diffusion processes. The discussion starts with a review of the reflection principle.

Suppose that price in interval t, $P_t(\tau)$, follows some continuous diffusion process whose innovations are symmetrically distributed about zero.[34] The conditional variance of the discrete price change $\varepsilon(\tau) = P_t(\tau) - P_t(0)$ within interval t is $\tau\sigma_t^2$. Let the normalized discrete cumulative

[34] Using methods and formulas that appear in Shepp (1979), results similar to those presented in this section can be obtained for processes with constant positive drift. These were not used in this paper because they would greatly complicate the numeric analysis, especially when stochastic volatility is modeled.

price change distribution be represented by $\Phi(\varepsilon)$. For notational simplicity, the interval subscript t henceforth will be dropped; τ will range between 0 and 1 within interval t; the initial price, $P_t(0)$, will be assumed to be equal to zero; and ε without any other annotation will represent the price innovation over the complete interval, $\varepsilon(1)$.

The probability that price is below some threshold X at the end of the interval ($\tau = 1$) is

$$\Pr\left(P(1) > X\right) = \Phi\left(X/\sigma\right). \tag{A.1}$$

The probability that price crosses X anytime in the interval is derived using the well-known reflection principle. The following derivation uses set notation to simplify the more difficult derivations that follow. Let

$$M = \text{ Max}\left(P(\tau)\right) \text{ for } 0 \le \tau \le 1 \tag{A.2}$$

and let S_X be the set of all price paths for which $M > X$ so that S_X is the set of all paths that cross X sometime within the interval. The passage probability, $\Pr(S_X)$, is the probability that some event in S_X is realized. To apply the reflection principle, note that for each path that terminates at $\tau = 1$ with $P(1) > X$, there was a time τ^* prior to 1 at which $P(\tau^*)$ was equal to X. At that time, the probability that the process ends above X is equal to the probability that the process ends below X because the subsequent diffusion has symmetric innovations with no drift. The reflection principle therefore implies that for every path that reaches X and terminates above X, there will be a second path that reaches X and terminates below X. The probability that the process crosses X at least once must therefore be twice the probability of terminating above X, or

$$\Pr\left(S_X\right) = \begin{bmatrix} 2\left(1 - \Phi\left(X/\sigma\right)\right) & \text{for } X > 0 \\ 1 & \text{for } X \le 0 \end{bmatrix}. \tag{A.3}$$

Some integrals in the dynamic trading problem objective functions must be evaluated for all paths that cross X. To evaluate these integrals, an expression for the density function of ε given $M > X$ must be derived. Let $\text{ppd}(\varepsilon|S_X)$ represent the partial probability density function of ε for all paths in S_X. The partial density is the non-normalized conditional density. The complete density function of ε is $\text{ppd}\left(\varepsilon|S_X\right) + \text{ppd}\left(\varepsilon|\sim S_X\right) = \text{ø}\left(\varepsilon/\sigma\right)/\sigma$ where $\sim S_X$ denotes the complement of S_X (the set of all paths not in S_X) and $\text{ø}(z)$ is the density function corresponding to $\Phi(z)$. The integral of $\text{ppd}(\varepsilon|S_X)$ over S_X is $\Pr(S_X)$.

The partial density, ppd(ε|S_X), is the union of the partial densities of all paths that terminate above X and all paths that cross X but terminate below X. The latter density is the reflection of the density of all paths that cross X. Specifically, the partial density is

$$\text{ppd}\left(\varepsilon \middle| S_X\right) = \begin{cases} \dfrac{\phi(\varepsilon/\sigma)}{\sigma} & \text{for } \varepsilon > X \text{ and } X > 0 \\ \dfrac{\phi((2X-\varepsilon)/\sigma)}{\sigma} & \text{for } \varepsilon < X \text{ and } X > 0 \\ \dfrac{\phi(\varepsilon/\sigma)}{\sigma} & \text{for } X < 0 \end{cases} \tag{A.4}$$

Let pE(ε|S_L) represent the partial expectation of ε over S_X. It is

$$\begin{aligned} \text{pE}\left(\varepsilon \middle| S_X\right) &= \int_{-\infty}^{\infty} \varepsilon \ \text{ppd}\left(\varepsilon \middle| S_X\right) d\varepsilon \\ &= \begin{cases} \int_X^{\infty} \varepsilon \dfrac{\phi(\varepsilon/\sigma)}{\sigma} d\varepsilon + \int_{-\infty}^{X} \varepsilon \dfrac{\phi((2X-\varepsilon)/\sigma)}{\sigma} d\varepsilon & \text{for } X > 0 \\ 0 & \text{for } X \le 0 \end{cases} \\ &= \begin{cases} \int_X^{\infty} \varepsilon \dfrac{\phi(\varepsilon/\sigma)}{\sigma} d\varepsilon + \int_X^{\infty} (2X - y) \dfrac{\phi(y/\sigma)}{\sigma} dy & \text{for } X > 0 \\ 0 & \text{for } X \le 0 \end{cases} \\ &= \begin{cases} X \ \text{Pr}\left(S_X\right) & \text{for } X > 0 \\ 0 & \text{for } X \le 0 \end{cases} \end{aligned} \tag{A.5}$$

B. RANDOM EXECUTION SET PROPERTIES

This subsection derives probabilities, partial probability density functions and partial expectations associated with the limit price execution process. These expressions are convolutions of the expressions derived above where the mixing distribution is the fill probability distribution. Set theory is used to improve the exposition and to ensure that limiting arguments are proper.

The derivation proceeds in four steps. First, probability statements about complementary price path sets are made. Statements are then made about random price path sets and about the union of discrete disjoint sets of random price paths. Finally, a limit is then taken to obtain the desired results.

First consider the properties of complementary price path sets: Let S_i and S_j respectively be the sets of all price paths for which $M > X_i$ and $M > X_j$. The intersection of these two sets,

S_{i_j}, is the set of all paths for which $X_j > M > X_i$ where $X_j > X_i$. The intersection set S_{i_j} is the complement of S_i in S_j because S_i is a subset of S_j. This decomposition implies the following three properties of interest to this study:

$$\Pr\left(S_{i-j}\right) = \Pr\left(S_i\right) - \Pr\left(S_j\right) \tag{A.6}$$

$$\text{ppd}\left(\varepsilon \middle| S_{i-j}\right) = \text{ppd}\left(\varepsilon \middle| S_i\right) - \text{ppd}\left(\varepsilon \middle| S_j\right) \tag{A.7}$$

$$\text{pE}\left(\varepsilon \middle| S_{i-j}\right) = \text{pE}\left(\varepsilon \middle| S_i\right) - \text{pE}\left(\varepsilon \middle| S_j\right). \tag{A.8}$$

Now let set $S_X^{F(X)}$ be a randomly chosen subset of the set S_X where $F(X)$ is the probability that a path in S_X appears in the subset $S_X^{F(X)}$. (The distribution F(X) will be the limit order fill probability distribution.) The results of interest are:

$$\Pr\left(S_X^{F(X)}\right) = F(X)\Pr\left(S_X\right) \tag{A.9}$$

$$\text{ppd}\left(\varepsilon \middle| S_X^{F(X)}\right) = F(X)\ \text{ppd}\left(\varepsilon \middle| S_X\right) \tag{A.10}$$

$$\text{pE}\left(\varepsilon \middle| S_X^{F(X)}\right) = F(X)\ \text{pE}\left(\varepsilon \middle| S_X\right) \tag{A.11}$$

Now consider properties of a union of disjoint sets: Let S_{Union} be the union of the three disjoint sets $S_{1-2}^{F(X_1)}$, $S_{2-3}^{F(X_2)}$, and $S_3^{F(X_3)}$ where $X_1 < X_2 < X_3$. Expressions (A.6) through (A.11) imply the following three properties for S_{Union}:

$$\begin{aligned}\Pr\left(S_{Union}\right) &= \Pr\left(S_{1-2}^{F(X_1)}\right) + \Pr\left(S_{2-3}^{F(X_2)}\right) + \Pr\left(S_3^{F(X_3)}\right)\\ &= F(X_1)\Pr\left(S_{1-2}\right) + F(X_2)\Pr\left(S_{2-3}\right) + F(X_3)\Pr\left(S_3\right)\\ &= F(X_1)\left[\Pr\left(S_1\right) - \Pr\left(S_2\right)\right] + F(X_2)\left[\Pr\left(S_2\right) - \Pr\left(S_3\right)\right] + F(X_3)\Pr\left(S_3\right)\\ &= \left[F(X_3) - F(X_2)\right]\Pr\left(S_3\right) + \left[F(X_2) - F(X_1)\right]\Pr\left(S_2\right) + F(X_1)\Pr\left(S_1\right)\end{aligned} \tag{A.12}$$

$$\text{ppd}\left(\varepsilon \middle| S_{Union}\right) = \left[F(X_3) - F(X_2)\right]\text{ppd}\left(\varepsilon \middle| S_3\right) + \left[F(X_2) - F(X_1)\right]\text{ppd}\left(\varepsilon \middle| S_2\right) + F(X_1)\text{ppd}\left(\varepsilon \middle| S_1\right) \tag{A.13}$$

$$\text{pE}\left(\varepsilon \middle| S_{Union}\right) = \left[F(X_3) - F(X_2)\right]\text{pE}\left(\varepsilon \middle| S_3\right) + \left[F(X_2) - F(X_1)\right]\text{pE}\left(\varepsilon \middle| S_2\right) + F(X_1)\text{pE}\left(\varepsilon \middle| S_1\right) \tag{A.14}$$

Finally, consider the continuous generalization of these three properties. Let S_{Exec} be the union of the set of all disjoint sets $\left\{S_{i-(i+1)}^{F(X_i)}\right\}$ for $X_i > X_o$ and let $X_{i_+1} = X_i + dX$ where dX is infinitesimal. Let $f(X)$ be the density function associated with the probability function $F(X)$. The following three properties are associated with this set:

$$\Pr\left(S_{Exec}\right)=F\left(X_o\right)\Pr\left(S_o\right)+\int_{X_o}^{\infty} f\left(X\right)\Pr\left(S_X\right)dX \tag{A.15}$$

$$\text{ppd}\left(\varepsilon\middle|S_{Exec}\right)=F\left(X_o\right)\ \text{ppd}\left(\varepsilon\middle|S_o\right)+\int_{X_o}^{\infty} f\left(X\right)\ \text{ppd}\left(\varepsilon\middle|S_X\right)dX \tag{A.16}$$

$$\text{pE}\left(\varepsilon\middle|S_{Exec}\right)=F\left(X_o\right)\ \text{pE}\left(\varepsilon\middle|S_o\right)+\int_{X_o}^{\infty} f\left(X\right)\ \text{pE}\left(\varepsilon\middle|S_X\right)dX \tag{A.17}$$

Substituting (A.3) through (A.5) into these three expressions yields

$$\Pr\left(S_{Exec}\right)=F\left(X_o\right)2\left(1-\Phi\left(X_o/\sigma\right)\right)+\begin{bmatrix}\int_{X_o}^{\infty} f\left(X\right)2\left(1-\Phi\left(X/\sigma\right)\right)dX & \text{for } X_o>0\\ \int_{X_o}^{0} f\left(X\right)dX+\int_{0}^{\infty} f\left(X\right)2\left(1-\Phi\left(L/\sigma\right)\right)dX & \text{for } X_o\le 0\end{bmatrix} \tag{A.18}$$

$$\text{ppd}\left(\varepsilon\middle|S_{Exec}\right)$$
$$=F\left(X_o\right)\begin{bmatrix}\dfrac{\phi\left(\varepsilon/\sigma\right)}{\sigma}\\ \dfrac{\phi\left(\left(2X_0-\varepsilon\right)/\sigma\right)}{\sigma}\\ \dfrac{\phi\left(\varepsilon/\sigma\right)}{\sigma}\end{bmatrix}+\int_{X_o}^{\infty} f\left(X\right)\begin{bmatrix}\dfrac{\phi\left(\varepsilon/\sigma\right)}{\sigma} & \text{for } \varepsilon>X \text{ and } X>0\\ \dfrac{\phi\left(\left(2X-\varepsilon\right)/\sigma\right)}{\sigma} & \text{for } \varepsilon\le X \text{ and } X>0\\ \dfrac{\phi\left(\varepsilon/\sigma\right)}{\sigma} & \text{for } X\le 0\end{bmatrix}dX \tag{A.19}$$

This expression can be simplified by examining the regions of the integral. In the first region, ε > X_o and $X_o > 0$ so that the integral reduces to

$$\begin{aligned}&\frac{\phi\left(\varepsilon/\sigma\right)}{\sigma}\int_{X_o}^{\varepsilon} f\left(X\right)dX+\int_{z}^{\infty} f\left(\varepsilon\right)\frac{\phi\left(\left(2X-\varepsilon\right)/\sigma\right)}{\sigma}dX\\ &=\frac{\phi\left(\varepsilon/\sigma\right)}{\sigma}\left[F\left(\varepsilon\right)-F\left(X_o\right)\right]+\int_{\varepsilon}^{\infty} f\left(\varepsilon\right)\frac{\phi\left(\left(2X-\varepsilon\right)/\sigma\right)}{\sigma}dX\end{aligned} \tag{A.20}$$

In the second region, ε ≤ X_o and $X_o > 0$. There the integral reduces to

$$\int_{\varepsilon}^{\infty} f\left(\varepsilon\right)\frac{\phi\left(\left(2X-\varepsilon\right)/\sigma\right)}{\sigma}dX \tag{A.21}$$

In the third region, $X_o \le 0$, so that the integral reduces to

$$\frac{\phi\left(\varepsilon/\sigma\right)}{\sigma}\int_{X_o}^{\infty} f\left(X\right)dX=\frac{\phi\left(\varepsilon/\sigma\right)}{\sigma}\left(1-F\left(X_o\right)\right) \tag{A.22}$$

Summarizing these results yields

$$\text{ppd}\left(\varepsilon \middle| S_{Exec}\right) = F(X_o)\begin{bmatrix} \dfrac{\phi(\varepsilon/\sigma)}{\sigma} & \text{for } \varepsilon > X_o \text{ and } X_o > 0 \\ \dfrac{\phi((2L_o-\varepsilon)/\sigma)}{\sigma} & \text{for } \varepsilon \le X_o \text{ and } X_o > 0 \\ \dfrac{\phi(\varepsilon/\sigma)}{\sigma} & \text{for } X_o \le 0 \end{bmatrix}$$

$$+\begin{bmatrix} \dfrac{\phi(\varepsilon/\sigma)}{\sigma}\left[F(\varepsilon)-F(X_o)\right]+\int_{\varepsilon}^{\infty} f(\varepsilon)\dfrac{\phi((2X-\varepsilon)/\sigma)}{\sigma}dX & \text{for } \varepsilon > X_o \text{ and } X_o > 0 \\ \int_{\varepsilon}^{\infty} f(\varepsilon)\dfrac{\phi((2X-\varepsilon)/\sigma)}{\sigma}dX & \text{for } \varepsilon \le X_o \text{ and } X_o > 0 \\ \dfrac{\phi(\varepsilon/\sigma)}{\sigma}\left(1-F(X_o)\right) & \text{for } X_o \le 0 \end{bmatrix} \tag{A.23}$$

The final expression is

$$\begin{aligned} \text{pE}\left(\varepsilon \middle| S_{Exec}\right) &= F(X_o)X_o \Pr\left(S_{X_o}\right)+\int_{X_o}^{\infty} f(X)\begin{bmatrix} X\Pr(S_X) & \text{for } X>0 \\ 0 & \text{for } X\le 0\end{bmatrix}dX \\ &= F(X_o)X_o 2\left(1-\Phi(X_o/\sigma)\right)+\begin{bmatrix} \int_{X_o}^{\infty} f(X)X\Pr(S_X)dX & \text{for } X_o>0 \\ \int_{0}^{\infty} f(X)X\Pr(S_X)dX & \text{for } X_o\le 0\end{bmatrix} \\ &= F(X_o)X_o 2\left(1-\Phi(X_o/\sigma)\right)+\int_{Max(X_o,0)}^{\infty} f(X)X2\left(1-\Phi(X/\sigma)\right)dX \end{aligned} \tag{A.24}$$

C. The Objective Function Integrals

The objective function integrals are considered in two sections. The first section considers the integrals in the uninformed liquidity trader problem and in the informed trader problem. These are considered together because they are quite similar. The second section considers the somewhat different integrals that appear in the value-motivated trader problem.

Integrals in the Liquidity Trader Problem and in the Informed Trader Problem

The various valuation functions specified for the uninformed liquidity trader problem and for the informed trader problem all can be expressed in the following form:

$$W_t(P_t) = \text{Max E}\begin{bmatrix} W_{t+1}(P_{t+1})+C_1 & \text{unfilled limit order} \\ L_t^* - \varepsilon_t + C_2 & \text{limit order filled} \\ -(1-\pi)\tfrac{1}{2}S + C_3 & \text{market order filled} \\ W_{t+1}(P_{t+1})+C_4 & \text{no order submitted} \end{bmatrix} \tag{A.25}$$

where C_1, C_2, C_3 and C_4 are various constants. (In some specifications, the price state variable and/or the value-innovation ε_t do not appear.) Four integrals must be computed to evaluate the expectations in this general expression. They are

1. the probability of execution of a limit order, Pr(*Exec*),
2. the partial expectation of ε_t given that a limit order executes,
3. the partial expectation of $W_{t+1}(P_{t+1})$ for limit order submissions that fail to execute, and
4. the unconditional expectation of $W_{t+1}(P_{t+1})$.

The first two integrals are computed directly from the formulas derived in the previous section. In those specifications where the value function does not depend on a price state variable, the third and fourth integrals are trivial: The third integral is simply $(1-\Pr(Exec))W_{t+1}$ and the fourth is just W_{t+1}. Otherwise, the third integral must be computed using the partial density function for ε_t given no execution, because the probability that a limit order executes in interval t depends on the realization of ε_t. The last integral is computed using unconditional densities.

To obtain expressions for these integrals, recall from Section II.C that limit orders are assumed to execute with probability $F(\delta)$ if at any τ in interval t,

$$L_t < P_t(\tau) - \delta . \tag{A.26}$$

This inequality is equivalent to

$$L_t^* + \tfrac{1}{2} S + \delta < P_t(\tau) - P_t \ . \tag{A.27}$$

where $L_t^* = P_t - A_t$. The form assumed for $F(\delta)$ is

$$F(\delta) = \begin{cases} 0 & \text{for } \delta < -\tfrac{1}{2} S \\ B(\delta) & \text{for } -\tfrac{1}{2} S \le \delta \le \tfrac{1}{2} S \\ 1 & \text{for } \delta > \tfrac{1}{2} S \end{cases} \tag{A.28}$$

where $B(\delta)$ is the beta distribution scaled to rise from 0 to 1 over $\left[-\tfrac{1}{2}S, \tfrac{1}{2}S\right]$. Let $f(\delta)$ be the corresponding density function for $F(\delta)$ so that

$$f(\delta) = \begin{cases} 0 & \text{for } \delta < -\tfrac{1}{2} S \\ B'(\delta) & \text{for } -\tfrac{1}{2} S \le \delta \le \tfrac{1}{2} S \\ 0 & \text{for } \delta > \tfrac{1}{2} S \end{cases} \ . \tag{A.29}$$

Let

$$M = \text{Max}\left(P_t(\tau) - P_t\right) \text{ for } 0 \le \tau \le 1 \tag{A.30}$$

and let S_X be the set of all price paths such $M > X$ where $X = L_t^* + \frac{1}{2}S + \delta$.

The probability of execution

For a given volatility, the probability that a limit order executes is given by

$$\Pr\left(Exec \middle| \sigma\right) = \int_{-\frac{1}{2}S}^{\frac{1}{2}S} \Pr\left(S_X\right) f\left(\delta\right) d\delta \ . \tag{A.31}$$

Transform the index of integration from δ to $X = L_t^* + \frac{1}{2}S + \delta$ to obtain

$$\Pr\left(Exec \middle| \sigma\right) = \int_{L_t^*}^{L_t^*+S} \Pr\left(S_X\right) f^*\left(X\right) dX \tag{A.32}$$

where

$$f^*(X) = f\left(X - L_t^* - \tfrac{1}{2}S\right) \ . \tag{A.33}$$

To apply (A.18), let the lower bound of integration $X_o = L_t^*$ since $f^*(X_o)$ is zero for any smaller value. The result is

$$\Pr\left(Exec \middle| \sigma\right) = \begin{cases} \int_{L_t^*}^{L_t^*+S} f^*\left(X\right) 2\left(1 - \Phi\left(X/\sigma\right)\right) dX & \text{for } L_t^* > 0 \\ \int_{L_t^*}^{0} f^*\left(X\right) dX + \int_{0}^{L_t^*+S} f^*\left(X\right) 2\left(1 - \Phi\left(L/\sigma\right)\right) dX & \text{for } L_t^* \le 0 \end{cases} \tag{A.34}$$

where $\sigma^2 = \sigma_{V,t}^2$. The infinite upper bound of integration that appears in (A.18) has been replaced by $L_t^* + S$ because the integral need only be evaluated up to the point where $f^*(X)$ is zero.

The unconditional probability of execution is obtained by integrating the conditional probability of execution over an assumed Inverted Gamma distribution for volatility:

$$\Pr\left(Exec\right) = \int_0^\infty \Pr\left(Exec \middle| \sigma\right) h\left(\sigma \middle| \overline{\sigma}, \nu\right) d\sigma \tag{A.35}$$

where $h(\sigma)$ is the Inverted Gamma density with an unconditional mean of $\overline{\sigma}$ and degrees of freedom ν.

The partial expectation of ε_t

Using (A.24), the partial expectation $\text{pE}(\varepsilon_t|Exec)$ for a given volatility is likewise computed as

$$\text{pE}\left(\varepsilon_t \middle| Exec, \sigma\right) = \int_{Max\left(L_t^*, 0\right)}^{L_t^*+S} f^*\left(X\right) X\, 2\left(1 - \Phi\left(X/\sigma\right)\right) dX \ . \tag{A.36}$$

The unconditional partial expectation is obtained by integrating over the Inverted Gamma volatility distribution.

The partial expectation of the value function

When a reservation price constraint is added to the informed traders' problem, the partial expectation of $W_{t+1}(P_{t+1})$ given no order execution must be computed. To compute this integral, the price process that the trader expects must be considered. It is

$$\begin{aligned} P_{t+1} &= V_{t+1} + e_{t+1} \\ &= V_t + \varepsilon_t + \theta_t + e_{t+1} \\ &= P_t - e_t + \varepsilon_t + \theta_t + e_{t+1} \\ &= P_t - \hat{e}_t - \psi_t + \varepsilon_t + \theta_t + \hat{e}_{t+1} + \psi_{t+1} \end{aligned} \tag{A.37}$$

where $\psi_t = e_t - \hat{e}_t$ is the trader's pricing error estimate error. The partial expectation must be taken over appropriate densities for ε_t, θ_t, ψ_t, and ψ_{t+1}. Fortunately, for a given volatility, the problem can be reduced to a bivariate integral by noting that limit order execution depends only on ε_t. The integral is computed over the partial density of the intra-interval value innovation ε_t, given no limit order execution, and then over the unconditional density of the sum of the other variables. Since these two densities are assumed to be independent, the integral can be evaluated sequentially. This analysis assumes that the density of the sum is given by

$$\theta_t + \psi_t + \psi_{t+1} \sim \mathrm{N}\left(\mu, k\sigma_t^2\right) \tag{A.38}$$

where k is some positive constant and σ_t^2 is the common stochastic volatility factor. Since the volatility factor is distributed with an Inverted Gamma distribution, the unconditional distribution of the sum is Student-t.

For a given volatility, the partial density function of ε_t given no execution can be expressed using (A.23). In particular,

$$\mathrm{ppd}\left(\varepsilon \middle| S_{NoExec}, \sigma\right) = \frac{\phi(\varepsilon/\sigma)}{\sigma} - \mathrm{ppd}\left(\varepsilon \middle| S_{Exec}, \sigma\right) \tag{A.39}$$

where

$$\text{ppd}\left(\varepsilon \middle| S_{Exec},\sigma\right)=\begin{cases}\dfrac{\phi(\varepsilon/\sigma)}{\sigma}F^{*}(\varepsilon)+\displaystyle\int_{\varepsilon}^{L_t^*+S} f^{*}(\varepsilon)\frac{\phi((2X-\varepsilon)/\sigma)}{\sigma}dX & \text{for } \varepsilon > L_t^* \text{ and } L_t^*>0 \\ \dfrac{\phi(\varepsilon/\sigma)}{\sigma} & \text{for } L_t^* \le 0\end{cases} \quad \text{(A.40)}$$

is simplified from (A.23) by recalling that $F^*(X_0)=F^*(L_t^*)=0$ and $f^*(X)=0$ for $X<L_t^*$. The unconditional partial density function expectation is obtained by integrating over the Inverted Gamma volatility distribution.

The partial expectation of the value function is computed over an approximation lattice of values for ε and for the sum $\theta_t+\psi_t+\psi_{t+1}$. The partial density ppd($\varepsilon|S_{NoExec}$) must be computed for each value of ε. The results are then used to integrate $W_{t+1}(P_t-\hat{e}_t-\psi_t+\varepsilon_t+\theta_t+\hat{e}_{t+1}+\psi_{t+1})$ over the lattice.

The unconditional expectation of the value function

The last integral is the unconditional expectation of $W_{t+1}(P_t-\hat{e}_t-\psi_t+\varepsilon_t+\theta_t+\hat{e}_{t+1}+\psi_{t+1})$. The bivariate integral is computed over an approximation lattice of values for ε_t and $\theta_t+\psi_t+\psi_{t+1}$. The unconditional densities are Student-*t* distributions implied by the assumed price processes.

Integrals in the Value-Motivated Trader's Problem

The objective function of the liquidity trader is

$$\begin{aligned} W(\hat{e}_t) &= \beta \mathrm{E}W(\hat{e}_{t+1}) + \text{Max E}\begin{cases}-K_L & \text{unfilled limit order}\\ L_t-V_{t+}-K_L-C_L & \text{limit order filled}\\ M_t-V_{t+}-K_M-C_M & \text{market order filled}\\ 0 & \text{no order submitted}\end{cases} \\ &= \beta \mathrm{E}W(\hat{e}_{t+1}) + \text{Max E}\begin{cases}-K_L\\ L_t^*-\varepsilon_t+\hat{e}_t+\tfrac{1}{2}S-K_L-C_L\\ -(1-\pi)\tfrac{1}{2}S+\hat{e}_t-K_M-C_M\\ 0\end{cases}\end{aligned} \quad \text{(A.41)}$$

Three integrals must be computed to evaluate the expectations in this expression. The first integral is the probability of execution of a limit order, Pr(*Exec*). The second integral is the partial expectation of ε_t given that a limit order executes. The last integral is the unconditional expectation of the value function, $\mathrm{E}W(\hat{e}_{t+1})$.

The probability of execution

The probability of execution is computed using a formula almost exactly the same as that derived for the liquidity trader problem. The only difference lies in the variance of the intra-interval diffusion process. In this problem, the variance σ^2 that appears in (A.34) is the variance of the sum of the value innovation process and the pricing error innovation process, $s_t = \varepsilon_t + n_t$, which is given by $\sigma_V^2 + \sigma_N^2 + 2\sigma_{VN}$.

The partial expectation of ε_t

The partial expectation $\text{pE}(\varepsilon_t|Exec)$ is a bivariate integral computed over the partial densities of ε_t and n_t given an execution. The integral is evaluated by computing the expectation of ε_t given $s_t = \varepsilon_t + n_t$ and then integrating over the partial density of s_t given an execution.

The expectation of given $s_t = \varepsilon_t + n_t$ is

$$\text{E}(\varepsilon_t | s_t) = \beta_\varepsilon s_t \tag{A.42}$$

where

$$\beta_\varepsilon = \frac{\sigma_V^2 + 2\sigma_{VN}}{\sigma_V^2 + \sigma_N^2 + 2\sigma_{VN}} \tag{A.43}$$

The partial expectation of this expression with respect to the partial density of execution for s_t is

$$\begin{aligned} \text{pE}(\text{E}(\varepsilon_t | s_t) | S_{Exec}) &= \text{pE}(\beta_\varepsilon s_t | S_{Exec}) \\ &= \beta_\varepsilon \, \text{pE}(s_t | S_{Exec}) \end{aligned} \tag{A.44}$$

The partial expectation $\text{pE}(s_t | S_{Exec})$ is computed using the same formula as that derived for the liquidity trader problem except that, in this problem, the variance that appears in (A.36) is the variance of the sum $s_t = \varepsilon_t + n_t$.

The unconditional expectation of the value function

The final integral is the unconditional expectation of $\text{E}W(\hat{e}_{t+1})$. To properly evaluate this integral, the process that the trader believes generates his pricing error estimates must be known. Given the assumptions made in Section II.B, the pricing error follows an AR(1) process:

$$e_{t+1} = \phi e_t + n_t \tag{A.45}$$

so that the expected pricing error estimate is

$$\hat{e}_{t+1} = \phi \hat{e}_t \,. \tag{A.46}$$

Assuming that the error in the estimate is uncorrelated with the pricing error innovation, the variance of this estimate is given by

$$\mathrm{Var}\left(\hat{e}_{t+1} - e_{t+1}\right) = \phi^2 \mathrm{Var}\left(\hat{e}_t - e_t\right) + \sigma_N^2 \;\; . \tag{A.47}$$

The pricing error estimate therefore has the same mean as the pricing error, but a larger variance.

This study computes $\mathrm{E}W\left(e_{t+1}\right)$ to approximate $\mathrm{E}W\left(\hat{e}_{t+1}\right)$. The error in the approximation will be small if the pricing error estimate is precise, or if the pricing error persistence parameter ϕ is small. In any event, it does not affect the optimal strategy, which does not depend on this term. The integral is computed by numerically integrating $\mathrm{E}W\left(\phi e_t + n_t\right)$ over the unconditional Student-*t* density for n_t.

APPENDIX B

LIMITATIONS DUE TO PRICE DISCRETENESS

The use of continuous processes to model discrete prices may bias the results for analyses of low price securities or for analyses of very short intervals. This appendix discusses why the bias arises and why it is not easily avoided.

The problem concerns the set of feasible limit order prices. Exchange minimum price variation regulations restrict this set to prices that are an integral multiple of 1/8 dollar. Unfortunately, tractable solutions cannot be obtained to the dynamic programming problems specified in this study when limit prices are required to be integral multiples of 1/8 dollar. Instead, this study requires that the limit price margin, $L_t^* = L_t - A_t$, be an integral multiple of 1/8 dollar. Since innovations in the ask A_t are continuous, the implied feasible set of limit prices changes from interval to interval depending on the change in A_t. To impose discreteness directly on L_t would require the introduction of some process to round a continuous latent bid to an observed discrete bid. The rounding error then would have to be included in the execution mechanism. Even if the rounding errors were independent through time, the resulting complexity would make the model intractable. Since the rounding errors are serially dependent, a completely proper representation of the process seems hopeless.[35]

The bias arises because the model assumptions imply that the probability of execution of a limit order depends only on the placement of the order relative to the beginning-of-interval bid. A simple extreme example illustrates a worst case manifestation of the problem. Suppose a sell limit order placed at 12.5 cents (one tick) above an initial bid would execute if price were to rise 6.25 cents. Now further suppose that price changes by exactly ± 5 cents per interval. If limit price discreteness were properly characterized in this model, an order placed in the first period and never canceled would execute if price ever rises ten cents above its starting point. When limit price discreteness is characterized by making the limit price margin L_t^* discrete, the order

[35] Harris (1990b) discusses the statistical properties of discrete pricing rounding errors.

will never execute no matter how often price increases. At the beginning of each interval, it always will be reset at 12.5 cents above the current bid.

An implicit assumption in this extreme example makes the problem appear worse than it is. The unrounded price presumably is known at the beginning of this example since we are told that the order will execute if price rises 6.25 cents. In practice, price is only known to the nearest 1/8 dollar. At the time the order is placed, the unrounded bid might be anywhere within a 12.5 cent region (assuming that the rounding error is uniformly distributed around the observed bid). If it is in the top 5 cents of the region, the order will execute in the first period with probability 1/2. The unconditional probability of execution therefore is not zero as suggested. It is the product of the probability of an increase times the probability that the unrounded price is in the top 5 cents of the rounding interval, which is $0.5 \cdot (5/12.5) = 0.2$.

In either event, the effect of this specification problem will be to make the derived optimal limit order strategies more aggressive when volatility within the interval is small relative to the tick. The orders must be more aggressive because the bias makes them less likely to execute.

Table 1

Annotated list of notation.

Class	Notation	Description
Time	t	Time subscript that indices the beginning of intervals
	$t+$	Time subscript that indices the end of intervals
	T	Time subscript of deadline interval
	τ	Time subscript that indices time within an interval
Prices	A_t	Ask
	B_t	Bid
	S	Spread
	P_t	Midspread price $P_t = (A_t - B_t) / 2$
	L_t	Limit price
	L_t^*	Limit price expressed relative to the ask $L_t^* = L_t - A_t$
	M_t	Market order execution price
	P_o	A prespecified price
	P^R	A reservation price
	d	The discrete tick (minimum price variation for limit orders)
Values	V_t	Underlying "true" value
	ε_t	Decision interval value innovation
	θ_t	Inter-interval value innovation
	$\sigma^2_{V,t}$	Decision interval value innovation volatility
	ν	Inverted gamma volatility process degrees of freedom parameter
	μ	Inter-interval value innovation mean
Pricing errors	e_t	Pricing error $e_t = P_t - V_t$
	$\hat{e}_t$	Pricing error estimate
	ϕ	Pricing error estimate persistence parameter
	$\sigma^2_{N,t}$	Pricing error innovation process volatility
	$\sigma_{VN,t}$	Value and pricing error innovation covariance
	ψ_t	Pricing error estimate error $\psi_t = e_t - \hat{e}_t$
Execution processes	π	Market order price improvement parameter
	δ	Degree of execution difficulty parameter/index
	$F(\delta)$	Fill probability function
Costs	K_L	Cost of submitting a limit order
	K_M	Cost of submitting a market order
	C_L	Cost of executing a limit order
	C_M	Cost of executing a market order
Valuation functions	$W_t(P_t)$	Dynamic programming valuation function
	BMP_t	Benchmark price
	β	Discount factor

Table 2

Summary of model assumptions.

Panel A: Time Assumptions

A1. Time is divided into a series of discrete intervals.
A2. Traders submit an order (if any) only at the beginning of an interval.
A3. Traders may submit orders only if the market is open.
A4. Prices and values change within intervals and trades take place within intervals.
A5. Between intervals, prices and values may change but trades do not take place.
A6. All unfilled orders expire at the end of each interval.

Panel B: Order Submission Assumptions

B1. Traders trade only a single prespecified quantity.
B2. Traders submit either a single market order or a single limit order at the beginning of each interval, or they abstain from submitting an order.
B3. Limit order prices are discrete. They must be a integer multiple of the minimum price increment, *d*.

Panel C: Common Trader Expectations about Prices and Values

C1. Value process within the interval *t*, $V_t(\tau)$, follows Brownian motion with no drift and constant innovation variance $\sigma^2_{V,t}$.
C2. Value innovations between intervals, θ_t, have mean μ, and variance $k\sigma^2_{V,t}$.
C3. Stochastic volatility $\sigma^2_{V,t}$ is distributed i.i.d. inverted gamma with mean $\overline{\sigma}^2_V$ and ν degrees of freedom.
C4. The value innovation within interval t, ε_t, is uncorrelated with the between interval innovation, θ_t.
C5. Bid and ask quotes at the beginning of interval t are given by $\begin{array}{l} B_t = P_t - \frac{1}{2}S \\ A_t = P_t + \frac{1}{2}S \end{array}$ where P_t is the beginning of interval value of the midquote price.
C6. The midquote price is the sum of value and a pricing error: $P_t(\tau) = V_t(\tau) + e_t(\tau)$.
C7. Orders have no effect on values and prices.

(Continued)

Table 2, Continued

Panel D: Defining Assumptions for the Various Stylized Traders

Assumption	Trader Type: Liquidity	Informed	Value-motivated
D1. Private Information	Uninformed.	Have a single piece of private information useful for predicting future prices.	Receive a perpetual flow of private information useful for forecasting future values and are they are able to model the process by which prices diverge from values.
D2. Deadline	Must beat deadline.	Need not trade.	Need not trade.
D3. Reservation price constraint	None.	May apply.	None.
D4. Trade frequency	Can trade only once.	Can trade only once.	Can trade perpetually, but no more than once each period.

Panel E: Trader Assumptions about the Pricing Error Process

Assumption	Liquidity	Informed	Value-motivated
E1. Pricing error process	Not specified	Not specified	$e_t(\tau) = e_t + n_t(\tau)$ in interval *t*, and $e_{t+1} = \phi e_{t+}$ between intervals *t* and *t*+1
E2. Pricing error innovations $n_t(\tau)$	Not specified	Not specified	Brownian motion, no drift, constant variance $\sigma^2_{N,t}$ proportional to $\sigma^2_{V,t}$.
E3. Cov(ε_t,n_t)	Not specified	Not specified	$\sigma_{VN,t}$ is proportional to $\sigma^2_{V,t}$.
E4. Pricing error estimate within interval	$\hat{e}_t(\tau) = 0$	$\hat{e}_t(\tau) = \hat{e}_t$	$\hat{e}_t(\tau) = \hat{e}_t$
E5. Expected pricing error estimate	$\hat{e}_{t+1} = 0$	$\hat{e}_{t+1} = \phi\hat{e}_t$	$\hat{e}_{t+1} = \phi\hat{e}_t$
E6. Pricing error estimate error $\psi_t = e_t - \hat{e}_t$	Not specified	Mean zero, independent	Mean zero, independent of everything.

(Continued)

Table 2, Continued

Panel F: Assumptions about Order Execution Expectations

F1. Market orders (and marketable limit orders) execute immediately when they are submitted.

F2. The expected execution price for a market sell order is $EM_t = B_t + \pi \tfrac{1}{2} S = P_t - (1-\pi)\tfrac{1}{2} S$ where B_t is the beginning of interval bid price, and the expected price improvement is some constant fraction π of one-half the bid/ask spread, S.

F3. If a limit order executes, the limit price is the execution price.

F4. The probability that a limit order will execute is given by

F5. $\Pr(Exec) = \int_{-\infty}^{\infty} \Pr(Exec|\delta) f(\delta) d\delta$

where $F(\delta) = \begin{cases} 0 & \text{for } \delta < -\tfrac{1}{2} S \\ G(\delta) & \text{for } -\tfrac{1}{2} S \le \delta \le \tfrac{1}{2} S \\ 1 & \text{for } \delta > \tfrac{1}{2} S \end{cases}$

and $G(\delta)$ is a Beta distribution function defined over the interval $\left[-\tfrac{1}{2} S, \tfrac{1}{2} S\right]$.

Panel G: Valuation Functions Assumptions

G1. Traders are expected net present value profit maximizers.

G2. There are fixed costs for submitting orders and for executing trades that vary by order type.

Table 3

Transformed dynamic program valuation functions for various objective function specifications for the informed trader optimal order submission strategy problem.

Specification	Benchmark Price	Dynamic Objective Function[a]	Terminal Condition
Sales Price Maximization	0	$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) + \mu - K_L \\ L_t - P_t + \hat{e}_t - \varepsilon_t - K_L - C_L \\ M_t - P_t + \hat{e}_t - K_M - C_M \\ W_{t+1}(P_{t+1}) + \mu \end{bmatrix}$	$W_{T+1}(P_{T+1}) = 0$
Value-Added Maximization	$\hat{V}_{T+1}$	Same as above	Same as above
Implementation Shortfall Minimization	P_o	Same as above	$W_{T+1}(P_{T+1}) = \hat{e}_{T+1}$
Transaction Cost Minimization, Post-Interval Price	P_{T+1}	Same as above	$W_{T+1}(P_{T+1}) = \hat{e}_{T+1}$
Transaction Cost Minimization, Post-Interval Price	P_{t+s}	$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L \\ L_t - P_t + \hat{e}_t - \varepsilon_t - s\mu - \hat{e}_{t+s} - K_L - C_L \\ M_t - P_t + \hat{e}_t - s\mu - \hat{e}_{t+s} - K_M - C_M \\ W_{t+1}(P_{t+1}) \end{bmatrix}$	$W_{T+1}(P_{T+1}) = 0$
Transaction Cost Minimization, Beginning-of-Interval Price	P_t	$W_t(P_t) = \text{Max E} \begin{bmatrix} W_{t+1}(P_{t+1}) - K_L \\ L_t - P_t - K_L - C_L \\ M_t - P_t - K_M - C_M \\ W_{t+1}(P_{t+1}) \end{bmatrix}$	Same as above

[a] The four lines in the dynamic objective function correspond to the following conditions:

Line 1: A limit order is submitted but does not execute.
Line 2: A limit order is submitted and executes.
Line 3: A market order is submitted and executes.
Line 4: No order is submitted.

Table 4

Solutions to the liquidity traders' problem for various bid-ask spreads and price change volatilities. The objective function is to maximize the expected difference between the net sales price and the prevailing bid at the time the decision to trade was made.

Bid/Ask Spread in Cents

	12.5					25.0					37.5				
Intervals Until	Price Change STD in Cents					Price Change STD in Cents					Price Change STD in Cents				
Deadline	5	10	20	40	80	5	10	20	40	80	5	10	20	40	80
Panel A: Optimal limit price position expressed in cents above the best offer.															
Deadline	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd
1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-12.5	-12.5	0.0	0.0	0.0
2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-12.5	0.0	0.0	0.0	0.0
3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	0.0	0.0	0.0	0.0	12.5	25.0
4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	0.0	25.0	37.5
5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	12.5	25.0	50.0
6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	12.5	25.0	50.0
7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	12.5	25.0	50.0
8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	12.5	25.0	50.0
9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	12.5	25.0	50.0
10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.5	12.5	12.5	0.0	0.0	12.5	25.0	50.0
Panel B: Probability of execution in current interval, in percent.															
Deadline	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
1	55.8	74.9	86.9	93.4	96.7	34.4	55.8	74.9	86.9	93.4	56.3	70.5	64.4	80.7	90.1
2	55.8	74.9	86.9	93.4	96.7	34.4	55.8	74.9	86.9	93.4	56.3	43.0	64.4	80.7	90.1
3	55.8	74.9	86.9	93.4	96.7	34.4	55.8	74.9	74.4	93.4	24.2	43.0	64.4	68.7	77.4
4	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	64.4	57.8	71.3
5	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	45.4	57.8	65.5
6	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	45.4	57.8	65.5
7	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	45.4	57.8	65.5
8	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	45.4	57.8	65.5
9	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	45.4	57.8	65.5
10	55.8	74.9	86.9	93.4	96.7	34.4	55.8	53.2	74.4	86.9	24.2	43.0	45.4	57.8	65.5
Panel C: Probability that execution ultimately takes place at the deadline with a market order, in percent.															
Deadline	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
1	44.2	25.1	13.1	6.6	3.3	65.6	44.2	25.1	13.1	6.6	43.7	29.5	35.6	19.3	9.9
2	19.6	6.3	1.7	0.4	0.1	43.0	19.6	6.3	1.7	0.4	19.1	16.8	12.6	3.7	1.0

3	8.7	1.6	0.2	0.0	0.0	28.2	8.7	1.6	0.4	0.0	14.5	9.6	4.5	1.2	0.2
4	3.8	0.4	0.0	0.0	0.0	18.5	3.8	0.7	0.1	0.0	11.0	5.5	1.6	0.5	0.1
5	1.7	0.1	0.0	0.0	0.0	12.1	1.7	0.3	0.0	0.0	8.3	3.1	0.9	0.2	0.0
6	0.7	0.0	0.0	0.0	0.0	8.0	0.7	0.2	0.0	0.0	6.3	1.8	0.5	0.1	0.0
7	0.3	0.0	0.0	0.0	0.0	5.2	0.3	0.1	0.0	0.0	4.8	1.0	0.3	0.0	0.0
8	0.1	0.0	0.0	0.0	0.0	3.4	0.1	0.0	0.0	0.0	3.6	0.6	0.1	0.0	0.0
9	0.1	0.0	0.0	0.0	0.0	2.2	0.1	0.0	0.0	0.0	2.7	0.3	0.1	0.0	0.0
10	0.0	0.0	0.0	0.0	0.0	1.5	0.0	0.0	0.0	0.0	2.1	0.2	0.0	0.0	0.0

Panel D: Sales price standard deviation, in cents per share.

Deadline	4.4	4.4	4.4	4.4	4.4	8.8	8.8	8.8	8.8	8.8	13.3	13.3	13.3	13.3	13.3
1	10.4	14.8	21.0	30.0	42.6	16.2	21.9	30.4	42.7	60.4	18.9	21.9	38.1	52.9	74.3
2	11.8	15.7	21.7	30.4	42.8	19.5	24.8	32.3	44.0	61.2	20.8	30.2	41.9	55.3	75.9
3	12.1	15.8	21.7	30.4	42.8	21.1	25.5	32.4	43.9	84.9	25.4	33.4	42.5	68.8	95.9
4	12.1	15.7	21.6	30.4	42.8	21.9	25.5	42.7	70.7	87.5	28.2	34.8	42.4	82.6	139.0
5	12.1	15.7	21.6	30.4	42.8	22.2	25.5	46.6	78.2	87.8	30.0	35.4	50.7	96.1	150.7
6	12.0	15.7	21.6	30.4	42.8	22.4	25.5	48.2	80.7	87.8	31.2	35.6	54.6	102.2	154.1
7	12.0	15.7	21.6	30.4	42.8	22.4	25.5	48.9	81.5	87.8	32.0	35.6	56.5	105.1	155.0
8	12.0	15.7	21.6	30.4	42.8	22.4	25.5	49.1	81.7	87.8	32.5	35.6	57.4	106.4	155.1
9	12.0	15.7	21.6	30.4	42.8	22.4	25.5	49.2	81.7	87.8	32.9	35.6	57.8	106.9	155.1
10	12.0	15.7	21.6	30.4	42.8	22.4	25.5	49.2	81.7	87.8	33.1	35.6	58.0	107.0	155.1

(Continued)

Table 4, *Continued*

Bid/Ask Spread in Cents

Intervals Until Deadline	12.5					25.0					37.5				
	Price Change STD in Cents					Price Change STD in Cents					Price Change STD in Cents				
	5	10	20	40	80	5	10	20	40	80	5	10	20	40	80
Panel E: Expected value of the dynamic programming problem, in cents per share.															
Deadline	0.1	0.1	0.1	0.1	0.1	3.3	3.3	3.3	3.3	3.3	6.4	6.4	6.4	6.4	6.4
1	2.5	2.8	2.9	3.0	3.0	7.0	8.1	8.7	9.0	9.2	13.3	13.6	14.2	14.9	15.2
2	3.5	3.4	3.3	3.2	3.1	9.5	10.3	10.1	9.8	9.6	16.4	17.0	17.0	16.6	16.1
3	4.0	3.6	3.3	3.2	3.1	11.1	11.2	10.5	9.9	9.6	18.5	19.0	18.0	16.9	16.2
4	4.2	3.6	3.3	3.2	3.1	12.2	11.7	10.5	9.9	9.6	20.1	20.1	18.4	17.0	16.2
5	4.2	3.6	3.3	3.2	3.1	12.9	11.8	10.6	9.9	9.6	21.4	20.7	18.5	17.0	16.2
6	4.3	3.7	3.3	3.2	3.1	13.3	11.9	10.6	9.9	9.6	22.3	21.1	18.6	17.1	16.3
7	4.3	3.7	3.3	3.2	3.1	13.6	12.0	10.6	9.9	9.6	23.0	21.3	18.7	17.1	16.3
8	4.3	3.7	3.3	3.2	3.1	13.8	12.0	10.6	9.9	9.6	23.6	21.4	18.7	17.1	16.3
9	4.3	3.7	3.3	3.2	3.1	14.0	12.0	10.6	9.9	9.6	24.0	21.5	18.7	17.1	16.3
10	4.3	3.7	3.3	3.2	3.1	14.1	12.0	10.6	9.9	9.6	24.3	21.5	18.7	17.1	16.3
Panel F: Expected number of intervals until execution.															
Deadline	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.4	0.3	0.1	0.1	0.0	0.7	0.4	0.3	0.1	0.1	0.4	0.3	0.4	0.2	0.1
2	0.6	0.3	0.1	0.1	0.0	1.1	0.6	0.3	0.1	0.1	0.6	0.7	0.5	0.2	0.1
3	0.7	0.3	0.2	0.1	0.0	1.4	0.7	0.3	0.3	0.1	1.2	1.0	0.5	0.4	0.3
4	0.8	0.3	0.2	0.1	0.0	1.6	0.8	0.6	0.3	0.1	1.7	1.1	0.5	0.6	0.4
5	0.8	0.3	0.2	0.1	0.0	1.7	0.8	0.8	0.3	0.1	2.0	1.2	0.8	0.7	0.5
6	0.8	0.3	0.2	0.1	0.0	1.8	0.8	0.8	0.3	0.2	2.3	1.3	1.0	0.7	0.5
7	0.8	0.3	0.2	0.1	0.0	1.8	0.8	0.9	0.3	0.2	2.5	1.3	1.1	0.7	0.5
8	0.8	0.3	0.2	0.1	0.0	1.8	0.8	0.9	0.3	0.2	2.7	1.3	1.1	0.7	0.5
9	0.8	0.3	0.2	0.1	0.0	1.9	0.8	0.9	0.3	0.2	2.8	1.3	1.2	0.7	0.5
10	0.8	0.3	0.2	0.1	0.0	1.9	0.8	0.9	0.3	0.2	2.9	1.3	1.2	0.7	0.5

The following parameter values are common to all solutions presented in this table:

Execution costs:	$C_L = K_L = 3$ cents per share
Order submission costs:	$K_L = 0.2$; $K_M = 0$ cents per share
Intra-interval price change drift	$\mu = 0$ cents per share/interval
Expected market order price improvement:	$\pi = 50\%$ of half-spread
Limit order fill probability distribution:	Uniform
Minimum price variation	$d = 12.5$ cents
Inverted gamma degrees of freedom	$\nu = 3$

Table 5

Solutions to the liquidity trader's problem for a baseline specification and for various changes in the baseline specification. The baseline specification is described at the bottom of the table.

		Change to Baseline Specification								
Intervals Until Deadline	**Baseline Solution**	Begin-of Interval Price Benchmark	Expected Mk Order Price Imp $\pi = 0$	Limit Fill Prob Distn Mean =0.25	Limit Fill Prob Distn Mean = 1	Mk Order Comm Cost C_M = 2¢	Limit Ord Subm Cost K_L = 0.4¢	One Cent Tick	Inter-Interval Drift μ = 0.5¢	Normal Innova-tions $\nu = \infty$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Optimal limit price position expressed in cents above the best offer.										
Deadline	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd
1	0.0	50.0	0.0	0.0	MkOrd	0.0	0.0	-1.0	0.0	0.0
2	0.0	75.0	0.0	0.0	MkOrd	0.0	0.0	0.0	0.0	0.0
3	12.5	100.0	0.0	0.0	MkOrd	12.5	0.0	8.0	NoSub	12.5
4	12.5	112.5	12.5	12.5	MkOrd	12.5	0.0	15.0	NoSub	12.5
5	12.5	125.0	12.5	12.5	MkOrd	12.5	0.0	17.0	NoSub	25.0
6	12.5	137.5	12.5	12.5	MkOrd	12.5	0.0	18.0	NoSub	25.0
7	12.5	150.0	12.5	12.5	MkOrd	12.5	0.0	18.0	NoSub	25.0
8	12.5	162.5	12.5	12.5	MkOrd	12.5	0.0	18.0	NoSub	25.0
9	12.5	162.5	12.5	12.5	MkOrd	12.5	0.0	18.0	NoSub	25.0
10	12.5	175.0	12.5	12.5	MkOrd	12.5	0.0	18.0	NoSub	25.0
Panel B: Probability of execution, in percent.										
Deadline	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
1	86.9	43.6	86.9	87.3	100.0	86.9	86.9	87.9	86.9	75.8
2	86.9	29.8	86.9	87.3	100.0	86.9	86.9	86.9	86.9	75.8
3	74.4	20.4	86.9	87.3	100.0	74.4	86.9	78.8	0.0	53.9
4	74.4	17.0	74.4	75.2	100.0	74.4	86.9	72.0	0.0	53.9
5	74.4	14.2	74.4	75.2	100.0	74.4	86.9	70.1	0.0	35.6
6	74.4	12.0	74.4	75.2	100.0	74.4	86.9	69.1	0.0	35.6
7	74.4	10.1	74.4	75.2	100.0	74.4	86.9	69.1	0.0	35.6
8	74.4	8.6	74.4	75.2	100.0	74.4	86.9	69.1	0.0	35.6
9	74.4	8.6	74.4	75.2	100.0	74.4	86.9	69.1	0.0	35.6
10	74.4	7.4	74.4	75.2	100.0	74.4	86.9	69.1	0.0	35.6
Panel C: Expected value of the dynamic programming problem, in cents per share.										
Deadline	3.3	3.3	-3.0	3.3	3.3	4.3	3.3	3.3	3.3	3.3
1	9.0	33.0	8.2	14.1	3.3	9.2	8.8	9.1	9.1	8.8

2	9.8	51.9	9.7	15.5	3.3	9.8	9.5	9.8	9.9	10.1
3	9.9	66.0	9.9	15.7	3.3	9.9	9.6	9.9	10.4	10.4
4	9.9	77.4	9.9	15.7	3.3	9.9	9.7	9.9	10.9	10.6
5	9.9	87.1	9.9	15.7	3.3	9.9	9.7	9.9	11.4	10.7
6	9.9	95.6	9.9	15.7	3.3	9.9	9.7	9.9	11.9	10.7
7	9.9	103.1	9.9	15.7	3.3	9.9	9.7	9.9	12.4	10.8
8	9.9	109.9	9.9	15.7	3.3	9.9	9.7	9.9	12.9	10.8
9	9.9	116.2	9.9	15.7	3.3	9.9	9.7	9.9	13.4	10.8
10	9.9	121.9	9.9	15.7	3.3	9.9	9.7	9.9	13.9	10.8

The following parameter values are used in the baseline specification:

Objective:	Maximize the expected difference between the net sale price and the prevailing bid at the time the decision to trade was made.
Bid/ask spread:	S = 25 cents per share
Price change STD:	σ_V = 40 cents per interval
Inter-interval drift:	μ = 0 cents per interval
Execution costs:	$C_L = K_L$ = 3 cents per share
Order submission costs:	K_L = 0.2; K_M = 0 cents per share
Expected market order price improvement:	π = 50% of half-spread
Limit order fill probability distribution:	Uniform (Mean = 0.5)
Minimum price variation	d = 12.5 cents
Inverted gamma degrees of freedom	ν = 3

Table 6

Solutions to the informed trader’s sell problem for a baseline specification and for various changes in the baseline specification. The baseline specification is described at the bottom of the table.

		Change to Baseline Specification				
Beginning Interval Intervals Benchmark Until Deadline	**Baseline Solution**	Initial Pricing Error $\hat{e}_{10}$ = 1000¢	Pricing Error Persistence ø = 0.5	Pricing Error Persistence ø = 0.9	End of Interval Price Benchmark	of Price
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Optimal limit price position expressed in cents above the best offer.						
Deadline	0.0	0.0	NoSub	MkOrd	NoSub	62.5
1	0.0	0.0	NoSub	0.0	NoSub	87.5
2	0.0	0.0	NoSub	0.0	NoSub	100.0
3	0.0	0.0	0.0	0.0	NoSub	112.5
4	0.0	-12.5	0.0	0.0	NoSub	125.0
5	0.0	-12.5	0.0	0.0	NoSub	137.5
6	0.0	-12.5	0.0	0.0	NoSub	150.0
7	-12.5	MkOrd	0.0	0.0	NoSub	162.5
8	-12.5	MkOrd	-12.5	-12.5	NoSub	175.0
9	-12.5	MkOrd	MkOrd	-12.5	NoSub	175.0
10	MkOrd	MkOrd	MkOrd	-12.5	NoSub	187.5
Panel B: Probability of execution, in percent.						
Deadline	86.9	86.9	0.0	100.0	0.0	36.0
1	86.9	86.9	0.0	86.9	0.0	24.6
2	86.9	86.9	0.0	86.9	0.0	20.4
3	86.9	86.9	86.9	86.9	0.0	17.0
4	86.9	96.7	86.9	86.9	0.0	14.2
5	86.9	96.7	86.9	86.9	0.0	12.0
6	86.9	96.7	86.9	86.9	0.0	10.1
7	96.7	100.0	86.9	86.9	0.0	8.6
8	96.7	100.0	96.7	96.7	0.0	7.4
9	96.7	100.0	100.0	96.7	0.0	7.4
10	100.0	100.0	100.0	96.7	0.0	6.3
Panel C: Expected value of the dynamic programming problem, in cents per share.						
Deadline	7.6	22.3	0.0	130.2	0.0	25.7
1	12.7	35.7	0.0	149.5	0.0	43.1
2	19.4	52.5	0.0	166.9	0.0	56.5
3	28.9	76.2	0.4	185.9	0.0	67.4
4	42.4	110.3	3.2	206.8	0.0	76.7
5	61.7	160.2	9.0	230.1	0.0	85.0
6	89.3	231.5	20.6	255.9	0.0	92.3
7	129.7	333.8	43.9	284.7	0.0	99.0
8	187.9	480.7	92.2	316.7	0.0	105.1
9	271.0	690.8	190.8	352.6	0.0	110.7
10	390.8	990.8	390.8	392.5	0.0	116.0

(Continued)

Table 6, Continued

Beginning Interval Intervals Benchmark Until Deadline	**Baseline Solution** (1)	Initial Pricing Error $\hat{e}_{10}$ = 1000¢ (2)	Pricing Error Persistence ø = 0.5 (3)	Pricing Error Persistence ø = 0.9 (4)	End of Interval Price Benchmark (5)	of Price (6)
		Change to Baseline Specification				
Panel D: Expected pricing error estimate, $\hat{e}_t$, in cents per share.						
Deadline	11.3	28.2	0.4	139.5	11.3	11.3
1	16.1	40.4	0.8	155.0	16.1	16.1
2	23.1	57.6	1.6	172.2	23.1	23.1
3	32.9	82.4	3.1	191.3	32.9	32.9
4	47.1	117.6	6.3	212.6	47.1	47.1
5	67.2	168.1	12.5	236.2	67.2	67.2
6	96.0	240.1	25.0	262.4	96.0	96.0
7	137.2	343.0	50.0	291.6	137.2	137.2
8	196.0	490.0	100.0	324.0	196.0	196.0
9	280.0	700.0	200.0	360.0	280.0	280.0
10	400.0	1000.0	400.0	400.0	400.0	400.0

The following parameter values are used in the baseline specification:

Objective:	Maximize expected value-added from trading
Bid/ask spread:	S = 25 cents per share
Value innovation STD:	σ_V = 40 cents per interval
Inter-interval drift:	μ = 0 cents per interval
Execution costs:	$C_L = K_L$ = 3 cents per share
Order submission costs:	K_L = 0.2; K_M = 0 cents per share
Expected market order price improvement:	$\pi\_$ = 50% of half-spread
Limit order fill probability distribution:	Uniform (Mean = 0.5)
Minimum price variation	d = 12.5 cents
Initial expected pricing error estimate	$\hat{e}_{10}$ = 400 cents per share
Pricing error persistence rate:	ø = 0.7
Inverted gamma degrees of freedom	ν = 3

Table 7

Solution to the informed trader’s sell problem for the baseline specification when the order submission strategy must satisfy a reservation price constraint. The current price slack and the number of periods until the deadline are the two state variables of this dynamic program. The current price slack is the difference between the quote midpoint and the reservation price at the beginning of the period. The current slack values presented in this table are subsets of the numeric lattices used to solve the problem. The subsets were chosen to minimize the space necessary to present the results. In each period, the slack value that is equal to the expected pricing error is indicated by an asterisk. The optimum limit order price is expressed as cents above the best offer. The baseline specification is described at the bottom of the table.

Price Slack $P_t - P^R$ (cents)	Optimum Order (cents)	Prob. Order Fills (%)	Program Value (cents)
At deadline $\hat{e}_0$ = **11.3¢**			
-138	NoSub	0.0	0.0
-127	125.0	14.2	1.1
-115	112.5	17.0	1.3
-104	100.0	20.4	1.7
-92	87.5	24.6	2.0
-75	75.0	29.8	2.5
-64	62.5	36.0	3.1
-52	50.0	43.6	3.8
-41	37.5	52.5	4.6
-29	25.0	62.8	5.5
-17	12.5	74.4	6.5
-12	0.0	86.9	7.6
6	0.0	86.9	7.6
11*	0.0	86.9	7.6
17+	0.0	86.9	7.6
1 period until deadline $\hat{e}_1$ = **16.1¢**			
-139	NoSub	0.0	1.5
-128	125.0	14.2	2.7
-116	112.5	17.0	3.2
-105	100.0	20.4	3.8
-93	87.5	24.6	4.5
-76	75.0	29.8	5.4
-64	62.5	36.0	6.3
-53	50.0	43.6	7.3
-41	37.5	52.5	8.4
-30	25.0	62.8	9.5
-13	12.5	74.4	10.9
-7	0.0	86.9	12.2
16*	0.0	86.9	12.3
114+	0.0	86.9	12.7
2 periods until deadline $\hat{e}_2$ = **23.1¢**			
-138	NoSub	0.0	3.3
-127	125.0	14.2	5.1
-115	112.5	17.0	5.8
-104	100.0	20.4	6.7
-92	87.5	24.6	7.7
-81	75.0	29.8	8.8
-63	62.5	36.0	10.3
-52	50.0	43.6	11.7
-40	37.5	52.5	13.2
-29	25.0	62.8	14.9
-17	12.5	74.4	16.7
-11	0.0	86.9	18.5
23*	0.0	86.9	18.8
167+	0.0	86.9	19.4
3 periods until deadline $\hat{e}_3$ = **32.9¢**			
-140	NoSub	0.0	5.4
-128	125.0	14.2	8.1
-117	112.5	17.0	9.2
-105	100.0	20.4	10.4
-88	87.5	24.6	12.2
-76	75.0	29.8	13.8
-65	62.5	36.0	15.5
-53	50.0	43.6	17.5
-42	37.5	52.5	19.8
-30	25.0	62.8	22.2
-13	12.5	74.4	24.9
-7	0.0	86.9	27.5
33*	0.0	86.9	28.0
234+	0.0	86.9	28.9
4 periods until deadline $\hat{e}_4$ = **47.1¢**			
-143	NoSub	0.0	7.9
-126	125.0	14.2	12.3
-114	112.5	17.0	13.9
-103	100.0	20.4	15.7
-91	87.5	24.6	17.7
-80	75.0	29.8	20.0
-68	62.5	36.0	22.6
-51	50.0	43.6	26.0
-39	37.5	52.5	29.1
-28	25.0	62.8	32.6
-16	12.5	74.4	36.4
-11	0.0	86.9	40.3
47*	0.0	86.9	41.3
260+	0.0	86.9	42.4
5 periods until deadline $\hat{e}_5$ = **67.2¢**			
-140	NoSub	0.0	11.5
-129	125.0	14.2	17.4
-117	112.5	17.0	19.6
-105	100.0	20.4	22.2
-88	87.5	24.6	25.8
-77	75.0	29.8	29.1
-65	62.5	36.0	32.8
-54	50.0	43.6	37.1
-42	37.5	52.5	41.8
-31	25.0	62.8	47.0
-13	12.5	74.4	52.9
-8	0.0	86.9	58.5
-2	0.0	86.9	58.7
4	-12.5	96.7	59.7
44	-12.5	96.7	60.0
50	0.0	86.9	60.0
67*	0.0	86.9	60.4
303+	0.0	86.9	61.7

(Continued)

Table 7, Continued

6 periods until deadline
$\hat{e}_6$ = 96.0¢

Current Slack $P_t - P^R$ (Cents)	Optimum Order (Cents)	Prob. Order Fills (%)	Program Value (Cents)
-140	NoSub	0.0	15.5
-128	125.0	14.2	24.0
-117	112.5	17.0	27.2
-105	100.0	20.4	30.9
-88	87.5	24.6	35.9
-77	75.0	29.8	40.7
-65	62.5	36.0	46.2
-54	50.0	43.6	52.4
-42	37.5	52.5	59.3
-31	25.0	62.8	67.0
-13	12.5	74.4	75.8
-2	0.0	86.9	84.5
4	-12.5	96.7	87.8
96*	-12.5	96.7	88.6
188	-12.5	96.7	88.9
194	0.0	86.9	88.9
280+	0.0	86.9	89.2

7 periods until deadline
$\hat{e}_7$ = 137.2¢

Current Slack $P_t - P^R$ (Cents)	Optimum Order (Cents)	Prob. Order Fills (%)	Program Value (Cents)
-93	87.5	24.6	47.8
-76	75.0	29.8	55.8
-64	62.5	36.0	63.8
-53	50.0	43.6	73.0
-41	37.5	52.5	83.3
-30	25.0	62.8	94.8
-18	12.5	74.4	107.4
-1	0.0	86.9	121.1
11	-12.5	96.7	127.8
16	MkOrd	100.0	127.9
22	-12.5	96.7	128.0
137*	-12.5	96.7	129.3
258+	-12.5	96.7	129.7

8 periods until deadline
$\hat{e}_8$ = 196¢

Current Slack $P_t - P^R$ (Cents)	Optimum Order (Cents)	Prob. Order Fills (%)	Program Value (Cents)
-17	12.5	74.4	152.0
-6	0.0	86.9	172.4
12	-12.5	96.7	184.8
18	MkOrd	100.0	186.8
104	MkOrd	100.0	186.8
110	-12.5	96.7	186.8
196*	-12.5	96.7	187.6
282+	-12.5	96.7	187.9

9 periods until deadline
$\hat{e}_9$ = 280¢

Current Slack $P_t - P^R$ (Cents)	Optimum Order (Cents)	Prob. Order Fills (%)	Program Value (Cents)
102	MkOrd	100.0	270.8
228	MkOrd	100.0	270.8
234	-12.5	96.7	270.8
280*	-12.5	96.7	270.9
343+	-12.5	96.7	271.1

10 periods until deadline
$\hat{e}_{10}$ = 400¢

Current Slack $P_t - P^R$ (Cents)	Optimum Order (Cents)	Prob. Order Fills (%)	Program Value (Cents)
262	MkOrd	100.0	390.8
400*	MkOrd	100.0	390.8
538	MkOrd	100.0	390.8

The following parameter values are used in the baseline specification:

Objective:	Maximize expected value-added from trading
Bid/ask spread:	S = 25 cents per share
Value innovation STD:	σ_V = 40 cents per interval
Nontrading interval STD:	σ_N = 1 cent per interval
Inter-interval drift:	μ = 0 cents per interval
Execution costs:	$C_L = K_L$ = 3 cents per share
Order submission costs:	K_L = 0.2; K_M = 0 cents per share
Expected market order price improvement:	π = 50% of half-spread
Limit order fill probability distribution:	Uniform (Mean = 0.5)
Minimum price variation	d = 12.5 cents
Initial expected pricing error:	e_{10} = 400 cents per share
Pricing error persistence:	ø = 0.7
Inverted gamma degrees of freedom	ν = 3

Table 8

Solutions to the value-motivated trader's sell problem for a baseline specification and for various changes in the baseline specification. The baseline specification is described at the bottom of the table.

Change to Baseline Specification

Estimated Pricing Error in Cents 0.1% to Quote Midpoint	**Baseline Solution** (1)	Pricing Error Innovation σ_N = 40¢ (2)	Pricing Error Persistence ø = 0.5 (3)	Value Innovation σ_V = 80¢ (4)	Correlation of n_t with ε_t 0.8 (5)	Correlation of n_t with ε_t -0.8 (6)	Discount Rate (1-ß)/ß = (6)
Panel A: Optimal limit price position expressed in cents above the best offer.							
-100	NoSub	325.0	NoSub	NoSub	500.0	NoSub	NoSub
-90	NoSub	300.0	NoSub	NoSub	450.0	NoSub	NoSub
-80	NoSub	275.0	NoSub	NoSub	412.5	NoSub	NoSub
-70	550.0	250.0	550.0	NoSub	362.5	NoSub	550.0
-60	475.0	225.0	475.0	NoSub	312.5	NoSub	475.0
-50	400.0	200.0	400.0	NoSub	275.0	NoSub	400.0
-40	325.0	175.0	325.0	NoSub	237.5	NoSub	325.0
-30	262.5	150.0	262.5	NoSub	200.0	NoSub	262.5
-20	187.5	125.0	187.5	600.0	162.5	NoSub	187.5
-10	125.0	100.0	125.0	375.0	125.0	NoSub	125.0
0	75.0	87.5	75.0	162.5	87.5	NoSub	75.0
10	25.0	62.5	25.0	25.0	62.5	0.0	25.0
20	0.0	50.0	0.0	0.0	37.5	-12.5	0.0
30	0.0	37.5	0.0	0.0	12.5	-12.5	0.0
40	0.0	25.0	0.0	0.0	0.0	-12.5	0.0
50	0.0	12.5	0.0	0.0	0.0	-12.5	0.0
60	-12.5	0.0	-12.5	0.0	0.0	MkOrd	-12.5
70	-12.5	0.0	-12.5	0.0	0.0	MkOrd	-12.5
80	-12.5	0.0	-12.5	0.0	0.0	MkOrd	-12.5
90	-12.5	0.0	-12.5	-12.5	-12.5	MkOrd	-12.5
100	-12.5	0.0	-12.5	-12.5	-12.5	MkOrd	-12.5
110	-12.5	-12.5	-12.5	-12.5	-12.5	MkOrd	-12.5
120	-12.5	-12.5	-12.5	-12.5	-12.5	MkOrd	-12.5
130	MkOrd	-12.5	MkOrd	-12.5	-12.5	MkOrd	MkOrd
140	MkOrd	-12.5	MkOrd	-12.5	-12.5	MkOrd	MkOrd
150	MkOrd	-12.5	MkOrd	-12.5	-12.5	MkOrd	MkOrd
160	MkOrd	-12.5	MkOrd	-12.5	-12.5	MkOrd	MkOrd
170	MkOrd	-12.5	MkOrd	-12.5	-12.5	MkOrd	MkOrd
180	MkOrd	-12.5	MkOrd	-12.5	MkOrd	MkOrd	MkOrd
190	MkOrd	-12.5	MkOrd	-12.5	MkOrd	MkOrd	MkOrd
200	MkOrd	MkOrd	MkOrd	-12.5	MkOrd	MkOrd	MkOrd
210	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd
220	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd
230	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd
240	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd
250	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd	MkOrd

(Continued)

Table 8, *Continued*

Estimated Pricing Error in Cents 0.1% to Quote Midpoint	**Baseline Solution** (1)	Pricing Error Innovation σ_N = 40¢ (2)	Pricing Error Persistence ø = 0.5 (3)	Value Innovation σ_V = 80¢ (4)	Correlation of n_t with ε_t: 0.8 (5)	Correlation of n_t with ε_t: -0.8 (6)	Discount Rate (1-ß)/ß = (6)

Change to Baseline Specification

Panel B: Probability of execution, in percent.

-100	0.0	4.1	0.0	0.0	1.4	0.0	0.0
-90	0.0	5.0	0.0	0.0	1.9	0.0	0.0
-80	0.0	6.1	0.0	0.0	2.3	0.0	0.0
-70	0.5	7.5	0.5	0.0	3.2	0.0	0.5
-60	0.8	9.4	0.8	0.0	4.7	0.0	0.8
-50	1.3	11.9	1.3	0.0	6.3	0.0	1.3
-40	2.2	15.2	2.2	0.0	8.6	0.0	2.2
-30	3.8	19.6	3.8	0.0	12.2	0.0	3.8
-20	8.2	25.6	8.2	2.3	17.6	0.0	8.2
-10	17.5	33.5	17.5	7.3	26.1	0.0	17.5
0	34.2	38.4	34.2	30.8	38.9	0.0	34.2
10	66.3	50.1	66.3	81.0	50.6	80.8	66.3
20	88.3	57.0	88.3	93.6	65.0	95.1	88.3
30	88.3	64.6	88.3	93.6	81.8	95.1	88.3
40	88.3	72.8	88.3	93.6	90.8	95.1	88.3
50	88.3	81.6	88.3	93.6	90.8	95.1	88.3
60	97.0	90.7	97.0	93.6	90.8	100.0	97.0
70	97.0	90.7	97.0	93.6	90.8	100.0	97.0
80	97.0	90.7	97.0	93.6	90.8	100.0	97.0
90	97.0	90.7	97.0	98.4	97.7	100.0	97.0
100	97.0	90.7	97.0	98.4	97.7	100.0	97.0
110	97.0	97.7	97.0	98.4	97.7	100.0	97.0
120	97.0	97.7	97.0	98.4	97.7	100.0	97.0
130	100.0	97.7	100.0	98.4	97.7	100.0	100.0
140	100.0	97.7	100.0	98.4	97.7	100.0	100.0
150	100.0	97.7	100.0	98.4	97.7	100.0	100.0
160	100.0	97.7	100.0	98.4	97.7	100.0	100.0
170	100.0	97.7	100.0	98.4	97.7	100.0	100.0
180	100.0	97.7	100.0	98.4	100.0	100.0	100.0
190	100.0	97.7	100.0	98.4	100.0	100.0	100.0
200	100.0	100.0	100.0	98.4	100.0	100.0	100.0
210	100.0	100.0	100.0	100.0	100.0	100.0	100.0
220	100.0	100.0	100.0	100.0	100.0	100.0	100.0
230	100.0	100.0	100.0	100.0	100.0	100.0	100.0
240	100.0	100.0	100.0	100.0	100.0	100.0	100.0
250	100.0	100.0	100.0	100.0	100.0	100.0	100.0

(Continued)

Table 8, *Continued*

Change to Baseline Specification

Estimated Pricing Error in Cents 0.1% to Quote Midpoint	**Baseline Solution** (1)	Pricing Error Innovation σ_N = 40¢ (2)	Pricing Error Persistence ø = 0.5 (3)	Value Innovation σ_V = 80¢ (4)	Correlation of n_t with ε_t 0.8 (5)	Correlation of n_t with ε_t -0.8 (6)	Discount Rate (1-ß)/ß = (6)
Panel C: Value of the dynamic programming problem, in dollars per share.							
-100	587.1	1286.3	453.5	578.5	647.6	539.3	293.1
-90	587.1	1286.4	453.5	578.5	647.7	539.4	293.2
-80	587.2	1286.4	453.5	578.6	647.7	539.4	293.2
-70	587.2	1286.4	453.5	578.6	647.8	539.4	293.2
-60	587.3	1286.5	453.6	578.7	647.8	539.5	293.3
-50	587.3	1286.5	453.6	578.7	647.9	539.5	293.3
-40	587.4	1286.6	453.6	578.8	647.9	539.6	293.4
-30	587.4	1286.7	453.7	578.8	648.0	539.6	293.5
-20	587.5	1286.8	453.7	578.9	648.1	539.7	293.5
-10	587.6	1286.8	453.8	579.0	648.2	539.8	293.6
0	587.7	1286.9	453.8	579.1	648.3	539.9	293.8
10	587.9	1287.0	453.9	579.3	648.5	540.0	293.9
20	588.1	1287.2	454.0	579.5	648.6	540.2	294.1
30	588.3	1287.3	454.2	579.7	648.8	540.4	294.3
40	588.5	1287.4	454.3	579.9	649.0	540.6	294.5
50	588.7	1287.6	454.5	580.1	649.2	540.8	294.7
60	588.9	1287.8	454.6	580.3	649.4	541.1	294.9
70	589.1	1288.0	454.8	580.6	649.7	541.3	295.2
80	589.4	1288.1	454.9	580.8	649.9	541.5	295.4
90	589.6	1288.3	455.1	581.1	650.1	541.8	295.6
100	589.9	1288.5	455.2	581.3	650.4	542.0	295.9
110	590.1	1288.7	455.4	581.6	650.6	542.3	296.1
120	590.4	1288.9	455.6	581.8	650.9	542.6	296.4
130	590.6	1289.1	455.7	582.1	651.1	542.8	296.7
140	590.9	1289.3	455.9	582.3	651.4	543.1	296.9
150	591.2	1289.5	456.1	582.6	651.6	543.4	297.2
160	591.4	1289.7	456.3	582.9	651.9	543.6	297.4
170	591.7	1289.9	456.4	583.2	652.2	543.9	297.7
180	592.0	1290.1	456.6	583.4	652.4	544.2	298.0
190	592.2	1290.3	456.8	583.7	652.7	544.5	298.3
200	592.5	1290.5	457.0	584.0	653.0	544.7	298.5
210	592.8	1290.7	457.1	584.2	653.3	545.0	298.8
220	593.1	1291.0	457.3	584.5	653.5	545.3	299.1
230	593.4	1291.2	457.5	584.8	653.8	545.6	299.4
240	593.6	1291.4	457.7	585.1	654.1	545.9	299.6
250	593.9	1291.6	457.9	585.4	654.4	546.1	299.9

The following parameter values are used in the baseline specification:

Objective:	Maximize expected present value-added from trading
Bid/ask spread:	S = 25 cents per share
Value innovation STD:	σ_V = 40 cents per interval
Inter-interval drift:	μ = 0 cents per interval
Execution costs:	$C_L = K_L$ = 3 cents per share
Order submission costs:	K_L = 0.2; K_M = 0 cents per share
Expected market order price improvement:	π = 50% of half-spread
Limit order fill probability distribution:	Uniform (Mean = 0.5)
Minimum price variation	d = 12.5 cents
Pricing error innovation STD:	σ_N = 20 cents per interval
Pricing error persistence:	ø = 0.7
Discount rate:	(1-ß)/ß = 0.05 percent per interval
Inverted Gamma degrees of freedom	ν = 3