Limit and Market Orders with Optimizing Traders

by

Praveen Kumar

College of Business Administration

University of Houston

and

Duane Seppi[1]

Graduate School of Industrial Administration

Carnegie Mellon University

September 5, 2000

[1]We thank Bruno Biais, Fischer Black, Peter Bossaerts, Joel Hasbrouck, Gur Huberman, Robert Jarrow, Ananth Madhavan, Vikram Nanda, Steve Shreve, Fallaw Sowell, Chester Spatt, Avanidhar Subrahmanyam, Chris Telmer, Raman Uppal, Mark Weinstein, Stan Zin and seminar participants at Alberta, UBC, Carnegie Mellon, Chicago, Columbia, Cornell, Minnesota, USC and the 1993 WFA meetings for helpful comments. We especially thank Rajdeep Singh for exemplary research assistance and helpful comments.

Abstract

We find equilibrium limit and market order placement strategies for optimizing liquidity traders, off–exchange liquidity providers called value traders and (possibly) a strategic informed trader. Nash (or Bayesian Nash) equilibria are found in two cases. The first has the exogenous arrival of information, but no informed trader. The second includes an informed trader which leads to endogenous information revelation and a "winner's curse" problem in limit order execution. Our main results are (a) the limit order book is random and is generated (along with market orders) by a linear factor model, (b) limit orders from value traders account for little of the depth in the limit book given even moderate brokerage costs, and (c) the limit order book is thinner in volatile than in "flat" markets. We also can construct a sufficient statistic for Bayesian updating with multiple order types when there is an informed trader.

1 Introduction

Investors trading on the New York Stock Exchange and other major exchanges have access to a rich set of order placement choices. In addition to the familiar market order, price–contingent limit orders specify a quantity to be bought (sold) if the price falls below (rises above) some prespecifed level. Stop orders allow investors to sell (buy) if the price falls below (rises above) a prespecified level. In practice price–contingent orders — particularly limit orders — play a prominent role in the transaction process. For example, over half of the submitted orders on the NYSE SuperDot electronic order submission system were limit orders in late 1993.[1]

Modeling the variety of order forms and how they are used by different types of investors represents one of the most important outstanding issues for microstructure theory. It is important for both practical and conceptual reasons. Investors naturally want to use these trading choices to their best advantage. Precisely how market and limit orders fit into an investor's optimal trading strategy depends, however, on both her particular motive for trade (e.g., private information, liquidity shocks, competitive liquidity provision) and on the interaction between her orders and the equilibrium trading decisions of other investors with potentially different trading motives. This suggest that investor heterogeneity is a prerequisite for a satisfactory theory of limit orders. A second reason limit orders are of interest is that the accumulated limit order book plays an important role both as a source of market liquidity for arriving orders and in the price formation process through which private information is aggregated and reflected in security prices. Furthermore, the availability of detailed order data makes it possible to begin empirically testing hypotheses about the composition of the limit (and stop) book.[2] Hence there is a need for models with price–contingent orders to generate these predictions.

Including limit orders in an equilibrium microstructure model of poses a number of conceptual and methodological challenges. This paper focuses on two key elements. The first is that all investors have symmetric access to both limit and market orders and their orders — including those of uninformed liquidity traders — are determined endogenously in an optimizing framework. In particular, we have

1. A continuum of liquidity traders who trade in response to liquidity shocks while wanting not

[1] See Chakravarty and Holden (1994). Harris and Hasbrouck (1992) report that 45 percent of the submitted orders for the 144 TORQ stocks in 1990/91 were limit orders.

[2] See Harris and Hasbrouck (1992) and Biais, Hillion and Spatt (1992).

to overpay (if buying) or to be underpaid (if selling).

2. Competitive uninformed off–exchange value traders who exploit opportunities in the limit order book to earn positive expected profits from liquidity provision.

3. A strategic informed investor who, if present, trades to exploit private prior knowledge about the future payoff of the asset.

This allows us to investigate behavioral interactions between different types of investors' strategies. For example, in contrast to the concentration of uninformed (market order) trading in Admati and Pfleiderer (1988), we find that uninformed trading by value traders (passive liquidity suppliers) "crowds out" orders form liquidity traders in parts of the limit order book. The second element in our model is that price discreteness, brokerage costs and other market imperfections are included and are shown to play an important role in the transaction process. Specifically,

1. A specialist sets prices to reflect the available information and also a non-informational bid/ask spread and then executes orders.

2. Traders incur up front brokerage costs to submit orders.

Both features represent extensions of existing models. The two cornerstones of the microstructure literature are the Kyle (1984, 1985) and Glosten and Milgrom (1985) models which restrict order submissions to market orders. Recently, however, Rock (1995), Glosten (1994) and Easley and O'Hara (1991) have presented models with price-contingent orders.[3] In Rock (1995) and Glosten (1994), limit and market orders are allowed while Easley and O'Hara (1991) allow stop and market orders. All three papers, however, restrict informed trading to market orders. The optimal strategy for strategic informed traders with limit and/or stop orders therefore remains unexplored. Rock and Glosten also restrict random liquidity trading to market orders as well. The result is a limit order book which is non–random and thus constant over time. Finally, none of these models include market imperfections.

In modeling the behavior of uninformed investors we consider two types of traders. As in Rock (1995), value traders do not need to trade per se, but they do submit orders if the expected profit is a non–negative. In constrast, our liquidity traders do actively need to trade. However, rather than

[3]In addition to these equilibrium models, Cohen, Maier, Schwartz and Whitcomb (1981), Angel (1992) and Harris (1994) derive optimal strategies for limit order placement given exogenous probabilities of execution.

treating them as "noise" traders with exogenous orders in Kyle (1985) we study optimizing liquidity traders who try to minimize the execution costs incured in responding to a rebalancing shock. Our work on this second group extends previous work by Spiegel and Subrahmanyam (1992), Seppi (1992) and Kumar and Seppi (1994) on optimal liquidity trading with market orders to a richer milieu with price–contingent orders.

Endogenizing the trading decisions of liquidity traders is important for two reasons. First, when investors can choose among combinations of market and limit orders, the "own" and the "cross" moments of the various market and limit order flows — which determine the execution probability of the limit orders — cannot be set arbitrarily. Rather, they should reflect a structure consistent with the optimization problems of other investors who submit orders. Thus, one does not want simply to assume a pattern of liquidity trading across the various types of orders. Second, the behaviour of the liquidity traders takes on additional significance in the presence of a strategic informed trader. Since an informed investor needs to hide in trading "noise" to exploit her information, the liquidity traders' trading strategy restricts the form of the informed's strategy in equilibrium. In particular, since the liquidity traders' orders are jointly determined by their optimization problem, any "noise" in the market order flows (due to liquidity trading) will be correlated with the "noise" in the limit order book. Thus, with $N+1$ possible market and limit orders, there will in general be fewer than $N+1$ sources of independent trading noise to "hide in." In fact, in our model — where aggregate liquidity order flows have an exact linear factor structure — there are only two sources of noise, an aggregate buying shock and an aggregate selling shock. Since our informed trader's orders must have the same factor structure (i.e., to avoid detection), she can be seen as simply picking her own personal shock realizations so as to maximize expected profits given her information.

The second element of our modeling approach is the inclusion of market imperfections — specifically, brokerage costs, price discreteness and a non–informational bid/ask spread. Empirically, the magnitudes of such imperfections (e.g., 3 to 10 cents a share brokerage fees or a 1/8th spread) are comparable to the estimated informational effects of trading. Hence, market imperfections might well be expected to have first–order effects on investors' trading strategies, market liquidity, and information revelation. Moreover, market imperfections are particularly important in models of limit orders. For example, it is precisely the non-informational bid/ask spread which opens the possibility of price improvement over market orders as a motive for limit orders (i.e., buying at the bid instead of the ask). In addition, transaction costs can have a significant impact on the depth and composition of the limit order book — specifically, the split between orders from liquidity traders (who are primarily interested in the relative cost of different orders) and the value traders (who only care about absolute costs).

We use our model to study two specific market environments. The first is a competitive market with value uncertainty due to an exogenous public announcement. The second adds an informed trader which leads to endogenous information revelation. In each market we calculate the equilibrium limit order book and market order flows, calculate summary statistics (e.g., expected costs and execution probabilities for limit orders) and investigate the effect of brokerage costs, exogenous price volatility, the probability of being informed and other factors.

Our main insights/contributions are:

1. Limit and market order flows have a linear factor structure in which one factor drives aggregate buying and another drives selling. This aggregate factor structure is inherited directly from the individual liqudity traders' optimization problem.

2. Our model features non–normal liquidity shocks. Without informed trading the shocks can have any distribution with the appropriate sign. With an informed trader exponentially distributed liquidity shocks lead to a tractable Bayesian updating problem.

3. Very low brokerage costs lead to a book dominated by non–random limit orders from the value traders as in Rock (1995). Higher costs produce a more random book driven by investors with other motives for trade (i.e., liquidity shocks or private information).

4. Informed trading leads to a "winner's curse" problem in limit order execution which curtails the use of limit orders by uninformed traders irrespective of any adverse selection effects in pricing. In particular, limit buys are more likely to be executed when the informed trader has bad news than good news.

The specialist's conditional expectation of the asset's value and its relation to the informed traders' strategy is a key methodological issue in models with private information. With multiple orders available for investors to use, this expectation depends in general on the entire limit order book as well as on the aggregate market orders. Rock (1995) and Glosten (1994) avoid this problem because their limit book is non-random and thus uninformative (i.e., since there is no "noise" for the informed trader to hide in). In our market the book is infomative. However, in our model the factor structure of orders leads to a sufficient "statistic" so that conditional expectations are still parsimonious.

In the interest of "truth in advertising," we disclose at the outset that this generality comes at a cost. First, the present analysis is limited to a static setting. We try, however, in a stylized way, to capture behavior which we view as important in dynamic markets. Third, our analysis of the informed trader market is restricted for tractability reasons to some simplifying parametric cases.

The paper proceeds as follows. Section 2 sets up the basic model in the absence of an informed trader, establishes that an equilibrium exists and presents a few numerical examples. Section 3 introduces an informed trader and then studies a special case of an endogenous valuation/winner's curse problem. Section 4 concludes with a discussion of robustness issues and possible extensions. All proofs are in the Appendix.

2 A Competitive Market with Exogenous Value Resolution

We begin by considering a one–period market in which three types of investors trade a risky stock.

The first are uninformed price–taking traders who either buy or sell for liquidity reasons. Specifically, there is a continuum of *liquidity traders* on $[0, \infty)$ who receive random rebalancing shocks (or targets) for that period. A buyer i_b receives a target to buy $w(i_b) > 0$ shares while a seller i_s wants to sell $w(i_s) < 0$. However, the orders they eventually submit are chosen to reflect a trade-off between expected trading costs and a quadratic penalty for deviations from their trading target.

The second type of traders are competitive risk–neutral *value traders* who exploit opportunities in the limit order book to earn positive expected profits. Being price takers, they collectively drive any expected profits there to zero. Unlike in Rock (1995), value trading does not preclude profit opportunities in the *realized* book since (a) our book is "closed" rather than "open" (i.e., investors cannot observe how many limit orders are already in the book) and (b) our book typically includes random orders from the liquidity traders.

The third trader is a competitive specialist who sets prices and then executes orders.

In this first model the stock's value changes due to a public announcement which arrives after orders are submitted, but *before* the specialist executes any trades. The probability distribution for the new value v is $h(v)$. This timing assumption is intended to capture the idea that in practice valuations can change between the submission and execution of limit orders. The resulting "value uncertainty" affects not only the prices at which investors trade (when they use market orders), but also which orders are actually executed (if limit orders are used). Section 3 adds an informed trader which leads to endogenous information revelation via the observed order flows.

Although value uncertainty and order flows are related emprically, the link is less than ironclad. Hasbrouck (1991) estimates that roughly 30 percent of the variance of permanent (informational) stock price changes are explained by order flow innovations. This leaves 70 percent of the variance to be explained by other factors such as informational spill–overs from other assets (e.g., related stocks, futures and other deriviative securities) and public announcements. Thus, our assumption of exogenous value realizations, while admittedly a simplification, is not an unreasonable polar case.[4]

The timing of events is in Figure 1. First, liquidity traders learn their individual targets $w(i)$. Then, given their beliefs about aggregate order submissions, they and the value traders each submit a profile of orders to the specialist for execution. The specialist observes v, sets his quotes given this information and executes orders.

Prices and values are related as follows. There is a finite grid $\mathbf{p}$ of prices at which trade occurs and on which limit orders can be posted. For concreteness we take $\mathbf{p} = \{p_1, p_2, p_3, p_4\}$. Interspersed in the price grid is a grid $\mathbf{v}$ of possible value realizations (or specialist expectations) v. Again for concreteness let $\mathbf{v} = \{v_1, v_2, v_3\}$ where

$$p_1 < v_1 < p_2 < v_2 < p_3 < v_3 < p_4. \tag{1}$$

Taking $v_2 = E(v)$, our grid allows traders to place ex ante "at the money" limit orders at p_2 and p_3 as well as "away from the money" orders at p_1 and p_4. This grid illustrates most of our economic insights and is readily generalized.

Transaction prices differ from the realized value v because of a non–informational bid/ask spread. The specialist's ask Q^a is the lowest price p above the realized value v and his bid Q^b is the highest price below v. This spread reflects either compensation for (unmodeled) variable costs incurred by competitive specialists or a form of specialist market power constrained by imperfect competition from floor brokers.

Investors may submit any combination of market and limit orders. *Market orders* are uncontingent orders to buy or sell a specified number of shares. Buy and sell orders are handled separately

[4]Rock (1995), Glosten (1994) and our model in Section 3 study the other polar case where all information is order related.

(rather than "batched" as in Kyle (1985)) because of the explicit bid/ask spread. Buy market orders trade at the ask and sell market orders at the bid. The specialist ensures that all market orders are executed.[5]

A non–informational bid/ask spread opens the possibility of price improvement on limit orders *vis–a–vis* market orders. A *buy limit order* indicates a willingness to buy shares at a specified price p_j on the grid $\mathbf{p}$. Intuitively, investors place limit buys because they hope to buy at the bid (given sufficient sellers on the sell side of the market) rather than at the ask (as with a market buy). The danger of using a limit buy order is either that (a) it might not be executed if not enough market sell orders arrive or (b) the value v may "run away" from the limit price p_j so that it is unexecuted (if prices rise) or is adversely "picked off" (if prices fall).

More formally, whether a limit buy order is filled or not depends on whether the specialist is willing (a) to better the limit price with a bid $Q^b > p_j$ (in which case it remains unexecuted), or (b) to match it with a bid $Q^b = p_j$ (in which case the specialist yields to public limit orders which are filled to the extent that there are sufficient sell orders at p_j) or (c) to take the other side of the limit buy himself by setting an ask $Q^a \leq p_j$ (in which case it is fully executed). In the event of partial execution due to insufficient sell orders the sell shares are allocated to limit buyers on the basis of "time priority" — which we model here in terms of random order arrival times. In particular, orders are submitted and assigned a position in the queue without investors knowing their priority (i.e., since the book is closed).

Similarly, a *sell limit order* indicates that the investor is willing to sell shares at a price p_j. Execution again depends on whether the specialist sets his ask $Q^a < p_j$ to better the limit price (no execution), sets his ask $Q^a = p_j$ to match it (partial execution depending on the number of buys) or sets his bid at or above $Q^b \geq p_j$ (full execution of "exposed" limit sells since the specialist is willing to take the other side).

Thus, the execution status of a limit order is a random variable depending on both (a) the realization v (which determines the location of the specialist's quotes Q^b and Q^a) and (b) the relative total supply and demand for shares at p_j — neither of which traders know when they submit their orders. The following notation represents these ideas. Let B_m denote the total shares submitted by *all* investors as market buy orders and B_j the total shares from limit buys at prices

[5]That is, in contrast to Rock (1990), the specialist is willing to absorb all executable orders which cannot be crossed with limit orders at his set quotes. Note also that Ginsi trading (i.e., trading different quantities at different prices to acheive an average price not on the price grid) is not optimal here. For example, selling some shares at a loss at the *bid*, and then the rest at the ask (after first filling all limit sells) simply reduces the specialist's profit.

p_j, $j = 1, \ldots, 4$. Similarly, let S_m denote total market sell orders and S_j the total shares submitted as limit sells at prices p_j. How much of a limit buy is filled is given by the "fill ratio"

$$F_{b,j} = \begin{cases} 1 & \text{with probability } \pi_{b,j} = \min\left[1, \frac{S_m+S_j}{B_j}\right] \quad \text{for } j = 1, \ldots, 4 \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

if ex post the specialist's bid matches the limit price p_j and is 0 or 1 otherwise as described above. The "min" reflects the fact that at most only 100 percent of the order can be filled. Similarly, the "fill ratio" for a limit sell is

$$F_{s,j} = \begin{cases} 1 & \text{with probability } \pi_{s,j} = \min\left[1, \frac{B_m+B_j}{S_j}\right] \quad \text{for } j = 1, \ldots, 4 \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

if the limit price p_j equals the specialist's ask and otherwise is 0 or 1 again.

The conditional "at the quote" fill ratios $F_{b,j}$ and $F_{s,j}$ (i.e., when the limit price equals the competing specialist quote) play an important role in our analysis. In particular, they describe limit order execution in states when execution is both profitable and possible (i.e., in that the specialist does not undercut/outbid them).

A brief review of the relevant rules on the NYSE may help to motivate our assumptions. Typically limit orders on the NYSE trade at the posted limit price. This is true even if the limit order is selling "cheap" (i.e., if p_j is below the bid) or buying "dear" (i.e., if p_j is above the ask). One exception is the "clean up price" rule 127 for block trades. Under this rule limit prices determine the priority of limit orders opposite a large block trade. If executed, however, they cross with the block at the block price, not the limit price. In this paper limit orders always trade at the posted limit price.[6]

We also assume that the specialist — rather than market orders on the other side of the market — is the counterparty for any exposed (i.e., cheap or dear) limit orders. That is, specialists,[7] by virtue of their advantaged position on the exchange floor, can interpose themselves between "exposed" limit orders and arriving market orders. Although awkward in a single period setting, in a dynamic market new information can easily arrive after limit orders are posted and before a market order happens to arrive. This assumption simplifies the model because market orders

[6]This is different from Kyle (???) where intramarginal limit orders are crossed at the marginal price.

[7]Strictly speaking, our specialist is a composite floor broker/specialist since NYSE specialists (as opposed to floor brokers) cannot trade against limit orders.

then trade at the realized bid and ask quotes rather than along a limit order schedule with random depth. In addition, intramarginal limit orders (i.e., above or below the realized quotes) do not first need to be subtracted off from the arriving market orders to find out how many "at the quote" limit orders can be crossed. Thus, while this assumption could be relaxed, it does simplify the analysis substantially.[8]

2.1 The individual liquidity trader's problem.

A generic liquidity trader i's objective reflects two considerations. He wants to minimize deviations of end–of–period shares traded from his target $w(i)$ while holding down expected trading costs from overpaying (on buy orders) and/or being underpaid (on sell orders). In particular, let $b_m(i)$ be the number of shares liquidity trader i submits as a market buy order and let $b_j(i)$ be the number of shares i posts as limit buy orders at prices p_j, $j = 1, \ldots, 4$. Similarly, let $s_m(i)$ and $s_j(i)$ denote the (unsigned) number of shares i submits as a market sell order and as limit sell orders at prices p_j respectively. The sign convention is that all order quantities are non-negative. For example, one cannot "buy" a negative number of shares (i.e., sell) at the ask via a market buy order. To sell, one must use sell orders. Vectors of these buy and sell orders are denoted by $\mathbf{b}$ and $\mathbf{s}$.

The expected trading costs per share on each of the possible market and limit buy orders are

$$\begin{aligned}
e_{b,1} &= h(v_1)(p_1 - v_1)E(\tilde{F}_{b,1}) + T \qquad (4)\\
e_{b,2} &= h(v_1)(p_2 - v_1) + h(v_2)(p_2 - v_2)E(\tilde{F}_{b,2}) + T\\
e_{b,3} &= h(v_1)(p_3 - v_1) + h(v_2)(p_3 - v_2) + h(v_3)(p_3 - v_3)E(\tilde{F}_{b,3}) + T\\
e_{b,4} &= h(v_1)(p_4 - v_1) + h(v_2)(p_4 - v_2) + h(v_3)(p_4 - v_3) + T\\
e_{b,m} &= h(v_1)(p_2 - v_1) + h(v_2)(p_3 - v_2) + h(v_3)(p_4 - v_3) + T.
\end{aligned}$$

These costs include both the implicit costs of over/underpaying relative to v as well as an up-front brokerage commission T per share. In particular, we assume T is paid on limit orders whether or not they are successfully executed. Similarly, the expected costs per share on sell orders are

$$e_{s,1} = h(v_1)(v_1 - p_1) + h(v_2)(v_2 - p_1) + h(v_3)(v_3 - p_1) + T \qquad (5)$$

[8]Market orders are less attractive if they trade at marginal (rather than fractionally at intramarginal) prices so under this assumption the limit order book should be deeper than otherwise.

$$
\begin{aligned}
e_{s,2} &= h(v_1)(v_1 - p_2)E(\tilde{F}_{s,2}) + h(v_2)(v_2 - p_2) + h(v_3)(v_3 - p_2) + T \\
e_{s,3} &= h(v_2)(v_2 - p_3)E(\tilde{F}_{s,3}) + h(v_3)(v_3 - p_3) + T \\
e_{s,4} &= h(v_3)(v_3 - p_4)E(\tilde{F}_{s,4}) + T \\
e_{s,m} &= h(v_1)(v_1 - p_1) + h(v_2)(v_2 - p_2) + h(v_3)(v_3 - p_3) + T.
\end{aligned}
$$

To foreshadow a bit, each order will be at best a "break even" proposition in equilibrium. The brokerage fee T and any over/underpaying relative to the true value v make trading costly. However, if T is not too big, limit orders can may still be profitable ex post for certain value realizations. For example, a limit buy at p_1 is profitable if executed given a value v_1. Thus, a limit buy at p_1 (or in general at any price p_j where the associated probability $h(v_j)$ is sufficiently high) breaks even ex ante only if the expected "at the quote" fill ratio $E(\tilde{F}_{b,1})$ (or $E(\tilde{F}_{b,j})$) is sufficiently low given T. This is precisely what the value traders ensure. A similar argument applies to sell limit orders.

With this notation liquidity trader i's problem is

$$
\begin{aligned}
\min_{\mathbf{b},\mathbf{s}} \quad & \left(\sum_{j=1}^{4} e_{b,j} b_j(i) + e_{b,m} b_m(i) + \sum_{j=1}^{4} e_{s,j} s_j(i) + e_{s,m} s_m(i) \right)^2 \qquad (6) \\
& + \phi \sum_{k=1}^{3} h(v_k) E\left[\left(b_k \tilde{F}_{b,k} + \sum_{l=k+1}^{4} b_l(i) + b_m(i) - \sum_{l=1}^{k} s_l(i) - s_{k+1}(i)\tilde{F}_{s,k+1} - s_m(i) - w(i) \right)^2 \right]
\end{aligned}
$$

subject to

$$
b_m(i), b_1(i), \ldots, b_4(i) \geq 0
$$

$$
s_m(i), s_1(i), \ldots, s_4(i) \geq 0.
$$

The quadratic specification has the virtue of tractability while capturing the intuition that limit orders involve a trade off between execution cost and probability.[9] In particular, the idea is that in a multiperiod market delayed execution is costly. In this single–period setting, however, the

[9] One aspect of the objective which deserves comment is the role of *squared* expected costs. If instead the objective was linear in the expected cost, the non-negativity constraint would force orders to be 0 for some w's and then affine in w everywhere else. The resulting piecewise solution would create a truncation problems for the aggregation of individual orders. Of course, this specification seems to suggest that liquidity traders also dislike *negative costs* (i.e., positive expected profits). However, the combination of up front brokerage costs and trading by the value traders ensure that costs are always non–positive in equlibrium Thus, far from manifesting a "kitchen sink" approach to realism, many of features of the model actually dovetail to improve tractability.

rebalancing penalty acts as a reduced form for any cost of delayed execution. The parameter ϕ allows us to vary the weighting of these two considerations with higher values of ϕ corresponding to greater dynamic impatience.

It is computationally helpful to identify and eliminate orders which a priori are never used (i.e., given any distribution $h(v)$ and beliefs about aggregate order flows).[10]

Proposition 1 *A liquidity buyer never uses market or limit sell orders or the (highest) limit buy order at p_4. A liquidity seller never uses market or limit buy orders or the (lowest) limit sell order at p_1.*

Given their quadratic objective liquidity traders' orders are linear in their individual shocks.

Proposition 2 *Optimal orders for a liquidity buyer i_b are linear in his realized shock $w(i_b)$*

$$\begin{aligned} b_m(i_b) &= \hat{\beta}_m(i_b)w(i_b) \\ b_j(i_b) &= \beta_j(i_b)w(i_b) \qquad j = 1, \ldots, 3 \end{aligned} \tag{7}$$

where the coefficients $\beta_m(i_b)$ and $\beta_j(i_b)$ are non-negative. Similarly, the optimal orders for a liquidity seller i_s also are linear in $w(i_s)$

$$\begin{aligned} s_m(i_s) &= \sigma_m(i_s)|w(i_s)| \\ s_j(i_s) &= \sigma_j(i_s)|w(i_s)| \qquad j = 2, \ldots, 4 \end{aligned} \tag{8}$$

where the coefficients $\sigma_m(i_s)$ and $\sigma_j(i_s)$ are again non-negative.

If liquidity traders are symmetrically informed about v and the S_j's and B_j's, then ϕ is the only individual–specific parameter on which the order coefficients $\beta_j(i)$ and $\sigma_j(i)$ depend. In particular, they do not not depend on the shock $w(i)$ (modulo its sign). We will assume that ϕ is identical for all liquidity traders.

[10] The symmetry of this result is a consequence of specifying the liquidity targets in terms of "shares bought or sold" rather than "cash expended or raised." If a cash target is used instead, then "cash raisers" still avoid buy orders, but "cash expenders" may simultanously use limit buy, market buy and also limit sell orders. To see why, consider a cash expender who submits a buy market order. High prices imply large cash outflows — potentially more than the cash target — while low prices may lead to underinvestment. This can be offset by using limit *sell* orders to reduce the net cash outflow at high prices and limit buy orders to augment it at low prices.

2.2 Value traders.

Value traders are simply risk neutral investors who trade if positive expected profits can be earned. Otherwise they have no particular need to trade. In Rock (1995) and Glosten (1994) the limit order book is composed entirely of orders from such traders.

Three differences between these models and ours deserve mention. First, value trading here is based on expectations of both random market and limit order flows whereas in Rock (1990) and Glosten (1992) limit orders are perfectly predictable and only market orders are random. Second, value assessments here change because of announcements and other exogenous reasons, whereas there (and in Section 3 here) values change only because of information revealed by order flows. Third, the off–exchange value traders in Rock (1990) and Glosten (1992) have a risk–bearing advantage over risk–averse specialists so that value trading is possible there despite the market makers' informational advantage. Here both value traders and specialists are risk neutral, but time priority together with price discreteness protects value traders against undercutting by the specialist.

The value traders are price takers. This eliminates strategic considerations as in Rock (1995). Being uninformed and free from liquidity shocks, value traders follow pure strategies so that their total orders, denoted here by $\alpha_{b,j}$ and $\alpha_{s,j}$ (or α in vector form), are perfectly predictable. Consequently, they are collectively modeled as submitting orders such that in equilibrium the expected costs (4) and (5) are non-negative for each order type. Since market orders always trade at a loss against the spread, value trading is confined to limit orders.

How does value trading affect expected costs? In (4) and (5) the only negative terms (i.e., ex post positive profits) are multiplied by the expected conditional fill ratios $E(\tilde{F}_{b,j})$ or $E(\tilde{F}_{s,j})$. Thus if the expected cost of any limit order were negative, value traders would submit such orders, thereby lowering the expected conditional fill ratios[11] and thus reducing the gain (in the profitable state) until any ex ante expected profits are eliminated.

The first step then is to show whether it is always possible to find non–negative orders $\alpha_{b,j}$ and $\alpha_{s,j}$ such that all of the expected costs in (4) and (5) are simultaneously non–negative. The next proposition gurantees that this is indeed possible.

Proposition 3 *There always exist non-negative value limit orders α such that the expected cost per share is simultaneously non-negative for all limit orders.*

[11]Substitute (9) and (10) below into (2) and (3).

The presence of value traders, in addition to its inherent naturalness, also justifies the squared cost minimization (rather than profit maximization) specification of the liquidity trader's objective. In equilibrium their trading ensures that expected costs are always non-negative.

2.3 Construction of aggregate order flows.

Although individually traders take the aggregate order flows as given, collectively their orders make up the total order flows. In aggregating their orders, the linearity of liquidity orders in the individual shocks $w(i)$ and the fact that the coefficients $\beta_m(i_b)$, $\beta_j(i_b)$, $\sigma_m(i_s)$, and $\sigma_j(i_s)$ are identical across buyers and sellers if beliefs and the impatience parameter ϕ are the same are very useful.

The aggregate orders in this market are

$$\begin{aligned} B_j &= \int_{i_b=0}^{\infty} b_j(i_b)di_b + \alpha_{b,j} \qquad j = m, 1, \ldots, 4 \\ &= \beta_j \int_{i_b=0}^{\infty} w(i_b)di_b + \alpha_{b,j} \end{aligned} \tag{9}$$

$$\begin{aligned} S_j &= \int_{i_s=0}^{\infty} s_j(i_s)di_s + \alpha_{s,j} \qquad j = m, 1, \ldots, 4. \\ &= \sigma_j \int_{i_s=0}^{\infty} |w(i_s)|di_s + \alpha_{s,j} \end{aligned} \tag{10}$$

Hence, for the aggregate order flows from liquidity buyers to be random we need to specify a process $\{w(i_b) : i_b \in [0, \infty)\}$ such that the integral giving the total liquidity buy shock

$$W_b = \int_{i_b=0}^{\infty} w(i_b)di_b \tag{11}$$

is random. A similar argument applies to aggregate liquidity sell orders.

Given the economics of this problem, the individual shocks $w(i_b)$ must have three properties. First, individual liquidity buyers must be "quantity small" (i.e., each trader i_b is a price–taker). Second, they must also be "informationally small" (i.e., each $w(i_b)$ is uninformative about W_b so that the conditional order flow moments used by i_b are simply the unconditional moments). Third, the strong law of large numbers is avoided when the individual shocks are aggregated. The first two conditions lead to a symmetric competitive market. The third ensures that the aggregate shock

W_b in (11) is random.[12]

One familiar process with this property is Brownian motion. Unfortunately, the resulting stochastic integral is normally distributed which is inappropriate since the aggregate shock W_b (i.e., total target shares to be purchased) must be non-negative. Another (also unsatisfactory) approach is to replicate the construction of Brownian motion by taking the limit of a process with independent, stationary, non–negative and infinitely divisible (but non–normal) increments.[13]. Breiman (1968) Proposition 14.19 gurantees the existence of such a process where the cummulative shock $W(i)$ through trader i is associated with the integral $\int_{j=0}^{i} w(j)dj$. However, by Breiman (1968) Proposition 14.20, the cummulative process $\{W(i) : i \in [0,1]\}$ has countably many jump discontinuities. Thus any trader i at a jump discontinuity has a shock $w(i)$ which is neither quantitity nor informationally small.

These difficulties arise because this approach tries to make each individual's shock $w(i)$ uninformative about the aggregate shock W by making it uninformative about the shock for *any* arbitrary subset of liquidity buyers.

An approach which does work is to partition a set of potential buyers $[0, \infty)$ into (a) a subset I_b of actual buyers with ex post identical shocks $w(i_b) = w_b^*$ and (b) a residual set of non-buyers for whom $w(i_b) = 0$. The common shock w_b^* can be any strictly positive constant or random variable. The aggregate shock is given then by the Lebesgue integral

$$\int_{i_b=0}^{\infty} w(i_b)di_b = w_b^* \mu(I_b) \tag{12}$$

where $\mu(I_b)$ is the Lebesgue measure of the set I_b. Although any individual buyer's shock ($i_b \in I_b$) is perfectly informative about other buyer's shocks, the desired independence is achieved by letting the measure $\mu(I_b)$ itself be a random variable given by $\mu(I_b) = \frac{W_b}{w_b^*}$. The numerator W_b may have any non–negative distribution with finite first and second moments (i.e., so that the individual's optimization probem (6) makes sense). Intuitively, aggregate uncertainty is driven not by *how much* individual traders buy, but rather by *how many* buy.

From (12) the buyers' shock $\int_{i_b=0}^{\infty} w(i_b)di_b$ is thus a well defined random variable equal to W_b.

[12]Aggregate randomness is particularly important in Section 3 since it means there is trading "noise" for the informed trader.

[13]An random variable is infinitely divisible if it has a distribution which for any integer $n > 0$ can be expressed as the distribution of the sum of n i.i.d. random variables. See Chung (1974, p. 239).

The sellers' aggregate shock,

$$W_s = \int_{i_s=0}^{\infty} |w(i_s)| di_s, \tag{13}$$

can be constructed analogously where the numerator W_s of the sellers' measure is taken to be independent of W_b.

This construction affords us great flexibility in specifying the liquidity shocks. This flexibility is useful in Section 3 where it allows us to choose a distribution which results in tractable Bayesian updating. Moreover, it is easily generalized to include multiple independent cohorts of buyers and sellers so that different investors can submit a variety of non-zero orders (rather than all the same order).

An immediate emprical implication of our aggregation results is that the total liquidity market and limit buy orders are generated by an exact one-factor model with a factor W_b and factor loadings equal to the individual liquidity order coefficients β_m, ..., β_4 from Proposition 2. Total liquidity sell orders are generated by another one-factor model with an independent factor W_s and factor loadings σ_m, ...,σ_4.[14] The possible factor representation of order flows is of interest for two reasons. First, as an empirical issue it would be possible to estimate how much of the accumulated limit order book (or alternately the limit order flows over an hour or some other horizon) is systematic/factor driven vs. idiosyncratic. Second, this factor structure has important implications for the informed trader's order placement strategy in Section 3.

2.4 Existence of equilibrium and examples.

This section first shows that a Nash equilibrium exists for the market we have described. We then compute and discuss some numerical examples.

The existence of a Nash equilibrium in this market hinges on finding a fixed point to the liquidity trader's optimization problem. Specifically, a Nash equilibrium is a set of liquidity order coefficients $(\beta, \sigma) = (\beta_m, \ldots, \beta_4, \sigma_m, \ldots, \sigma_4)$ such that, given the aggregate liquidity order flows $\beta_m W_b$, ..., $\beta_4 W_b$, $\sigma_m W_s$, ..., $\sigma_4 W_s$ and the value order flows $\alpha_{b,j}$ and $\alpha_{s,j}$ they induce, the coefficients (β, σ) are optimal responses for a generic liquidity buyer i_b and seller i_s. Such a fixed point does exists so that we have

Proposition 4 *A Nash equilibrium (β, σ) always exists in this market.*

[14]More generally, if investors are divided into M groups with different weightings ϕ then order flows will have an M-factor structure.

The proof involves first finding a fixed point in the first and second non-central moments of the six fill ratios $\tilde{F}_{b,j}$ and $\tilde{F}_{s,j}$ and then substituting these back into the liquidity trader's problem (6) to get the coefficient fixed point. The advantage of this trick is that the realized fill ratios — and hence their moments — lie in the interval $[0, 1]$ whereas the liquidity order coefficients (β, σ) are hard to bound a priori because of the fractional execution of "at the quote" limit orders.

Closed-form solutions are, unfortunately, not available, but equilibria are easy to calculate numerically. Table 1 provides some examples for a market with pricing on an "eights" grid $\mathbf{p} = \{20, 20\ \frac{1}{8}, 20\ \frac{1}{4}, 20\ \frac{3}{8}\}$ and an offset value grid $\mathbf{v} = \{20\ \frac{1}{16}, 20\ \frac{3}{16}, 20\ \frac{5}{16}\}$. The liquidity buy shock W_b (and likewise W_s) is exponentially distributed with[15]

$$Prob(W_b = W) = \begin{cases} \theta e^{-\theta W} & \text{if } W \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{14}$$

The probability density of an exponential aggregate shock is strictly decreasing in shock size which is not unreasonable. The parameter θ in these examples is picked so that the mean and standard deviation of the rebalancing shocks are $1/\theta = 500$ shares.

Market equilibria are calculated for (a) a variety of brokerage costs T ranging from a high of \$.05 (i.e., five cents a share) to a low of \$.001, (b) two impatience parameter values $\phi = 0.1$ and 1.0 and (c) three mean–preserving value distributions $h(v)$. For each equilibrium we report the ex ante expected costs $e_{b,1}$ and $e_{b,2}$ for limit buys at the two relevant prices p_1 and p_2,[16] the expected "at the quote" fill ratios $E(\tilde{F}_{b,1})$ and $E(\tilde{F}_{b,2})$, the optimal liquidity buyer's coefficients β_1, β_2 and β_m and the value traders' limit buys $\alpha_{b,1}$ and $\alpha_{b,2}$. Since these parameterizations are symmetric, the corresponding limit sells at p_4 and p_3 and market sells are symmetric with $\sigma_4 = \beta_1, \sigma_3 = \beta_2$ and $\sigma_m = \beta_m$.

The limit order books look reasonable. Not surprisingly, value investors place more limit orders close to the unconditional expected value at p_2 (for buys) and at p_3 (for sells) in "flat" markets (i.e., when the probability $h(v_2)$ is high) than in volatile markets. In Panel B, for example, value

[15] Since the exponential is a special case of the gamma distribution, we can, if desired, use the additivity of gamma variables to construct the aggregate shock (which is Gamma $(\theta, 1)$) as the sum of N independent Gamma $(\theta, 1/N)$ cohort shocks.

[16] Since the market here is symmetric, a limit buy (sell) at p_3 (p_2) is weakly dominated by a market order both for rebalancing purposes (since it may not be fully filled) and in terms of cost (i.e., from (4) and (5) it is only the ex post profitable orders which are not fully filled). Hence, $\beta_3 = \sigma_2 = 0$ so that limit buys are always posted at lower prices than limit sells. Note also that the expected cost for a market buy $e_{b,m}$ is simply half a "tick" or \$.0625 plus the brokerage cost T.

traders have no limit buys at p_2 given just a 40 percent chance of v_2, but when this probability is 90 percent their orders total 1842 shares. Liquidity traders' limit orders at p_2, however, are non-monotone (see Panels C and D). Increasing the liquidity traders' impatience parameter ϕ leads to greater use of market orders for liquidity trading and thus more aggresive limit order placement by value traders (see Panels B and D).

One interesting feature of these examples is the sensitivity of investors' trading strategies to brokerage costs. Even modest brokerage costs have a dramatic impact on the composition of orders in the limit book. In Panel A, for example, a cost of just 5 cents drives away all value trading, leaving only liquidity trading. In Panel C, however, dropping the cost to just 0.1 cents leads to a dramatically deeper book in which orders from the value traders swamp those of the liquidity traders (i.e., whose expected orders at p_2 total only about $.10 \times E(\tilde{W}_b)$ or 50 shares). Costs of 1 and 5 cent per share also lead to books concentrated at the ex ante "at the money" prices p_2 and p_3 rather than further "away from the money" at p_1 and p_4 — consistent with the evidence on order frequency in Harris and Hasbrouck (1992).

These examples suggest that limit order books on the NYSE and elsewhere may — due to brokerage and other trading costs — be determined primarily by investors who actively want to trade for liquidity or (in the next section) informational reasons rather than by disinterested liquidity providers as in Rock (1990). As an empirical matter, these books are qualitatively different since liquidity trading leads to a random limit order book whereas value trading leads to a deterministic book.[17] Thus in principle it should be possible to assess the relative importance of these two types of traders in practice.

The difference in their propensity to use limit orders arises because value traders — who do not *need* to trade per se — care only about the absolute dollar costs (or profits) of trading. In contrast, liquidity traders care about the relative cost of different ways of actively rebalancing. This insight not only explains comparative statics across parameterizations, but also explains why, for any particular parameterization, value and liquidity traders use different combinations of orders. A nice example of this is the first market in Panel B where value traders post limit buys only at p_1 (where expected trading profits just offset brokerage costs) while liquidity buyers mainly post limit orders at p_2 (where the expected cost of 0.5 cents precludes value orders but is still a bargain compared to 7.25 cents for market orders and where the likelihood of execution is still high).

[17] Of course, if the market parameterization itself is stochastic, then unconditionally the book is also random.

3 Equilibrium with a Strategic Informed Trader

Informed trading affects markets (and uninformed traders in particular) in three ways. First, order flows reveal information which complicates the specialist's pricing problem. Second, if sufficiently intense this information revelation affects not only the prices at which trade occurs, but also — unlike in models with just market orders — which orders are executed (i.e., which limit orders). Third, uninformed traders lose to informed traders not only via the usual price effects, but also through a "winner's curse" problem in limit order execution. Specifically, fewer limit buys are executed when the informed trader has good news than when she has bad news and conversely for limit sells.

The second effect, the impact of order flows on which limit orders are executed, is the focus of Rock (1995) and Glosten (1994), but is not yet tractable in our model. However, given price discreteness, there is a wide range of parameterizations for which this particular issue does not arise. For such parameterizations, we can still study the winner's curse problem.

The model again has a single round of trading. The timing of events is in Figure 2. At the beginning a strategic, risk-neutral *informed trader* observes an informative signal ψ. After seeing ψ, she trades to maximize her expected profits. Liquidity traders again learn their targets $w(i)$ and then, together with the value traders, submit orders for execution.

The risky security has an exogenous terminal payoff v^* with a binomial distribution

$$v^* = \begin{cases} \overline{v} & \text{with probability } \overline{h} \\ \underline{v} & \text{with probability } \underline{h} = 1 - \overline{h}. \end{cases} \tag{15}$$

The realized payoff is announced publicly only after trades are executed, but the informed investor may learn v^* beforehand with positive probability (see Figure 3). Conditional on good news $\bar{v}$, her signal ψ either reveals this (denoted ψ_u) with probability π or not (denoted ψ_n). Similarly, conditional on bad news $\underline{v}$, a signal ψ either reveals this (denoted ψ_d) with probability π or not (denoted ψ_n). A signal ψ_n is thus entirely uninformative.

We must distinguish between the exogenous terminal payoff v^* and the specialist's interim conditional expectation of v^* denoted by v. The interim expectation is an endogenous variable based on the now informative order flow. It is v which determines at what prices trade occurs and which orders are executed. In particular, transaction prices are again restricted to a discrete grid $\mathbf{p}$ with the specialist's bid Q^b assumed to be the next price below v and his ask Q^a the next price above v on the grid.[18]

Notation for liquidity and value orders is the same as in Section 2. For the informed trader, let $x_{b,m}$ and $x_{s,m}$ denote buy and sell market orders and $x_{b,j}$ and $x_{s,j}$ buy and sell limit orders at p_j. The total order flows B_m, B_j, S_m and S_j are redefined here to include the informed trader's orders and $\mathbf{B}$ and $\mathbf{S}$ again denote vectors of the total buy and sell orders respectively.

A Bayesian Nash equilibrium is a quintuple ($\mathbf{b}$, $\mathbf{s}$, α, $\mathbf{x}$, $\mathbf{v}$) such that investors follow optimal strategies where

1. Liquidity orders $\mathbf{b}$ and $\mathbf{s}$ are optimal given each trader i's target $w(i)$.

2. Total value orders α insure that the expected cost is non-negative for each type of order.

3. The informed's orders $\mathbf{x}$ maximize her expected profit.

and such that the specialist's interim valuation v is the conditional expectation of v^* given these strategies and the observed order flows (i.e., $v = E(v^* : \mathbf{B}, \mathbf{S})$).

The analysis becomes tractable with four additional assumptions.

1. There is an upper bound $\bar{g}$ on an appropriately scaled version of the informed's orders.

2. The specialist calls "trading halts" in response to non–factor orders.

3. Aggregate liquidity shocks are exponentially distributed.

4. The probability of being informed π is sufficiently small.

These assumptions are discussed in more detail below. Briefly, however, they restrict the support of the specialist's interim valuation v to at most four possible values. Thus, the (endogenous) interim value grid is $\mathbf{v} = \{v_1, v_2, v_3, v_4\}$. Assumption 1 prevents the informed trader from submitting unbounded orders (i.e., given the nature of Bayesian updating with exponential shocks under Assumption 3). It can be viewed as a reduced–form for risk aversion or other portfolio restrictions on the informed. Assumption 4 says that by picking π sufficiently low, the value grid $\mathbf{v}$ can be compressed within a single interval in the price grid $\mathbf{p}$. Although trading still reveals information (i.e., the specialist's posterior differs from his prior), the price impact is now less than a "tick." In

[18]To highlight the endogenous order flow effects, we assume that there are no post–submission/pre–execution announcements. Public announcements are easily added by treating the terminal payoff v^* as the sum of a payoff which is revealed by ψ plus a second independent payoff which is announced (as in Section 2) before orders are executed and which simply displaces the valuation induced by the first term either up or down.

this special case, the bid and ask quotes and thus the transaction prices themselves are unaffected by it. Hence, in the absence of public announcements, the price grid reduces to $\mathbf{p} = \{p_1, p_2\}$.

The independence of prices from order flows simplifies the analysis by reducing the informed's trading to a "plunging" strategy — that is, buying the maximal amount (under Assumption 1) given good news ψ_u and selling the maximal amount given bad news ψ_d. It does not, however, reduce informed trading to a matter of indifference for other traders. The winner's curse problem in limit order execution means liquidity and value traders are still hurt by informed trading. Thus their optimal trading strategies respond to it.

The analysis is divided into three steps. First, the optimal trading strategy for individual liquidity traders is found and then aggregate to give the total liquidity order flows. Second, the order placements by value traders and the informed investor are characterized given the liquidity traders' strategy. And third, we show that a fixed point exists in the liquidity trader strategy.

As in Section 2, a liquidity trader i's problem is to minimize squared expected trading costs plus the expectation of a quadratic rebalancing penalty. The choice variables are market buys and sells, limit buys at p_1 and limit sells at p_2. Note that a limit buy (sell) at p_2 (p_1) is indistinguishable from a market buy (sell) in this setting and thus can be ignored. Thus given a target $w(i)$, trader i's problem is

$$\begin{aligned} \min_{\mathbf{b},\mathbf{s}} \quad & (e_{b,1}b_1(i) + e_{b,m}b_m(i) + e_{s,2}s_2(i) + e_{s,m}s_m(i))^2 \\ & +\phi \sum_{v^* \in \{\underline{v},\bar{v}\}} h(v^*) E[(b_1\tilde{F}_{b,1} + b_m(i) - s_2(i)\tilde{F}_{s,2} - s_m(i) - w(i))^2] \end{aligned} \tag{16}$$

subject to

$$b_m(i), b_1(i) \geq 0$$

$$s_m(i), s_2(i) \geq 0$$

where the expected costs (per share) on limit and market orders are

$$\begin{aligned} e_{b,1} &= E[(p_1 - \tilde{v}^*)\tilde{F}_{b,1}] + T \\ e_{b,m} &= p_2 - E(\tilde{v}^*) + T \\ e_{s,2} &= E[(\tilde{v}^* - p_2)\tilde{F}_{s,2}] + T \\ e_{s,m} &= E(\tilde{v}^*) - p_1 + T \end{aligned} \tag{17}$$

and where $\tilde{F}_{b,1}$ and $\tilde{F}_{s,2}$ are again the "at the quote" fill ratios. Because of the winner's curse (see equations (25) and (26) below), the fill ratios and the terminal payoff v^* will be correlated.

Using the same logic as in Propositions 1 and 2, the optimal orders are again unique and linear in a buyer i_b's target $w(i_b)$

$$\begin{aligned} b_m(i_b) &= \beta_m(i_b)w(i_b) \\ b_1(i_b) &= \beta_1(i_b)w(i_b) \end{aligned} \tag{18}$$

and in a liquidity seller i_s's target $w(i_s)$

$$\begin{aligned} s_m(i_s) &= \sigma_m(i_s)|w(i_s)| \\ s_2(i_s) &= \sigma_2(i_s)|w(i_s)| \end{aligned} \tag{19}$$

where the coefficients $\beta_m(i_b)$, $\beta_1(i_b)$, $\sigma_m(i_s)$ and $\sigma_2(i_s)$ are non-negative. If liquidity traders are again uninformed price takers with identical beliefs and preference weights ϕ, these coefficients are identical across traders.

Individual orders are aggregated into total liquidity orders as before. Specializing the model to the case of i.i.d. exponentially distributed liquidity shocks W_b and W_s (as in the market examples in Section 2) leads to a particularly tractable form for the specialist's conditional expecation v.[19]

Value traders again fully exploit any opportunities to place orders which earn expected trading profits. The aggregate limit orders $\alpha_{b,1}$ and $\alpha_{s,2}$ simply insure that the expected costs $e_{b,1}$ and $e_{s,2}$ in (17) are non-negative.

This brings us to the main task of this section — the derivation of the informed investor's trading strategy and a (hopefully) parsimonious updating rule for the specialist. What makes this tractable is the insight that the order flows from liquidity traders are generated by an exact factor model. Consequently, the four order types here (i.e., market buy and sell and the two limit orders) do not constitute four independent sources of trading noise for the informed trader to "hide in," but implicitly only two — namely, the aggregate liquidity buying and selling shocks W_b and W_s. Hence, to pool with liquidity trading, the informed investor's orders must have the same factor structure

$$(x_{b,1}, x_{b,m}, x_{s,2}, x_{s,m}) = (\beta_1 g_b, \beta_m g_b, \sigma_2 g_s, \sigma_m g_s). \tag{20}$$

[19]Normality is not a possibility since buy and sell orders are submitted and aggregated separately.

The only difference is that she picks her personal factor realizations (g_b, g_s) to maximize expected profits given her signal ψ.

Our task, then, is to verify that under Assumptions 1 through 4 the informed investor in fact *chooses* to submit "factor orders" as in (20) and then to solve for (g_b, g_s) as a function of the signal ψ.

To say whether pooling with the liquidity traders is optimal for the informed trader we must first say what happens if she does not pool. Towards this end, first note that with factor liquidity orders and predictable value orders, non–factor deviations are perfectly detectible. To see this, consider the following normalization of the observed orders

$$\begin{aligned} A_{b,1} &= \tfrac{B_1-\alpha_{b,1}}{\beta_1} &= W_b + \frac{x_{b,1}}{\beta_1} \\ A_{b,m} &= \tfrac{B_m}{\beta_m} &= W_b + \frac{x_{b,m}}{\beta_m} \\ A_{s,2} &= \tfrac{S_2-\alpha_{s,2}}{\sigma_2} &= W_s + \frac{x_{s,2}}{\sigma_2} \\ A_{s,m} &= \tfrac{S_m}{\sigma_m} &= W_s + \frac{x_{s,m}}{\sigma_m}. \end{aligned} \tag{21}$$

The difference $A_{b,1} - A_{b,m} = \frac{x_{b,1}}{\beta_1} - \frac{x_{b,m}}{\beta_m}$ thus permits the specialist to measure exactly any difference between normalized informed buy orders. With factor orders this difference is zero. Analogous statements are true for $A_{s,2} - A_{s,m}$ and informed sell orders.

Non–factor orders lead to the following inference problem. Given, for example, non–factor buy orders $\frac{x_{b,1}}{\beta_1} \neq \frac{x_{b,m}}{\beta_m}$, should good news or bad news be infered? There are two ways to deal with this problem.

The first is to construct an equilibrium in which we (a) restrict attention to π's such that, given factor informed orders, the interim valuation v is always between p_1 and p_2 and likewise (b) specify beliefs such that "off equilibrium" prices (i.e., given non–factor informed orders) are also between p_1 and p_2. If this is possible, prices are independent of the informed's orders so that her optimal strategy is to take the maximum position (long or short) permitted under Assumption 1. The trick then is to specify the bound the informed's trading so that her maximal position itself satisfies the factor condition. Thus, in choosing to use maximal orders, the informed trader also (indirectly) chooses to use factor orders (20). Under this approach then the bound $\bar{g}$ in Assumption 1 is on the informed trader's normalized orders.

Bound on Informed Trading. *If $\beta_j \neq 0$, then $\frac{x_{b,j}}{\beta_j} \leq \bar{g}$ for $j = m, 1$ and likewise, if $\sigma_j \neq 0$, then*

$\frac{x_{s,j}}{\sigma_j} \leq \bar{g}$ *for* $j = m, 2$.

Showing that "in equilibrium" the interim valuation v stays between p_1 and p_2 is an issue we take up below. However, an "off equilibrium" valuation v between p_1 and p_2 is easily supported by beliefs that any non–factor order is equally likely to come from an informed trader with good news ψ_u as from one with bad news ψ_d. In this case, the interim valuation is unchanged from the unconditional expectation. More generally, any (non–constant) "off equilibrium" schedule taking values in the interval $[p_1, p_2]$ will work (i.e., good and bad news need not be equally likely for every non–factor order realization). Non–degenerate posteriors are not unreasonable since, for example, non–factor sell orders could be submitted in conjunction with factor buying (given good news) or not (given bad news).

A second approach is simply to allow our specialist, much like specialists on the NYSE, to call a "trading halt" when abnormal (non–factor) order flows are observed (i.e., order flows inconsistent with liquidity trading and hence which unambiguously reveal the presence of an informed trader). Clearly this also forces the informed trader to use factor orders.[20] As in the first approach, we again assume that, given factor orders, the (discrete) bid and ask prices do not depend on the realized order flows. This again leads to the use of the maximal feasible[21] factor orders.

This second approach squarely facing the main problem here. Markets with adverse selection problems require uninformed "noise" to function. Thus, precisely as in Kyle (1985) where the market is open for market orders only given noise in the market order flow, our market is open (i.e., ignoring the perfectly predictable value orders) only for factor orders given that there is only factor noise.

No matter how non–factor orders are excluded, the informed trader's strategy has two important properties. First, it has a factor structure (either because of off equilibrium beliefs or a trading halt rule). Second, it involves "plunging" or taking maximal positions given good or bad news (i.e., provided the interim valuation v is indeed always in the interval $[p_1, p_2]$) and otherwise sitting out.

[20]Although somewhat heavy handed, the specialist is unlikely to use this power to force other "unreasonable" equilbria on the market (e.g., ones in which no trading is ever allowed) since he in fact earns expected profits by keeping the exchange open.

[21]There may be some slack for some orders relative to their bounds if, under Assumption 1, these bounds are imposed on unnormalized rather than normalized orders. However, at least one bound will be binding.

Thus we have

$$(x_{b,1}, x_{b,m}, x_{s,2}, x_{s,m}) = \begin{cases} (\beta_1 \bar{g}, \beta_m \bar{g}, 0, 0) & \text{if } \psi_u \\ (0,0,0,0) & \text{if } \psi_n \\ (0,0,\sigma_2 \bar{g}, \sigma_m \bar{g}) & \text{if } \psi_d. \end{cases} \tag{22}$$

The particular form of the informed trader's strategy has two important consequences for Bayesian inference on the equilibrium path. First, because she only uses factor orders the (common) normalized aggregate order flows

$$A_{b,1} = A_{b,m} = A_b \tag{23}$$

$$A_{s,2} = A_{s,m} = A_s.$$

constitute a sufficient statistic for *all* order flows. The availability of a lower dimensional sufficient statistic greatly simplifies the specialist's conditional expectation v which otherwise depends on each of the different order flows. We should be clear that this observation about a sufficient statistic is a general property of markets in which both the informed and liquidity traders uses factor orders. In particular, it is not restricted to markets satisfying the simplifying assumptions made on π here.

Second, the "plunging" aspect of the informed trader's strategy, together with exponential liquidity shocks (Assumption 2), leads to a particularly simple form for the specialist's conditional expectations v. In terms of the four regions in (A_b, A_s) space defined by the bound $\bar{g}$ (see Figure 4) we have

$$v = \begin{cases} \frac{h(\bar{v})(1-\pi)}{0+h(\bar{v})(1-\pi)+(1-h(\bar{v})(1-\pi)+e^{\theta\bar{g}}h(\bar{v})\pi}(\bar{v}-\underline{v})+\underline{v} & \text{if } (A_b, A_s) \text{ is in region I} \\ \frac{h(\bar{v})(1-\pi)}{0+h(\bar{v})(1-\pi)+(1-h(\bar{v})(1-\pi)+0}(\bar{v}-\underline{v})+\underline{v} & \text{if } (A_b, A_s) \text{ is in region II} \\ \frac{e^{\theta\bar{g}}h(\bar{v})\pi+h(\bar{v})(1-\pi)}{e^{\theta\bar{g}}h(\bar{v})\pi+h(\bar{v})(1-\pi)+(1-h(\bar{v})(1-\pi)+e^{\theta\bar{g}}(1-h(\bar{v})\pi}(\bar{v}-\underline{v})+\underline{v} & \text{if } (A_b, A_s) \text{ is in region III} \\ \frac{e^{\theta\bar{g}}h(\bar{v})\pi+h(\bar{v})(1-\pi)}{e^{\theta\bar{g}}h(\bar{v})\pi+h(\bar{v})(1-\pi)+(1-h(\bar{v})(1-\pi)+0}(\bar{v}-\underline{v})+\underline{v} & \text{if } (A_b, A_s) \text{ is in region IV.} \end{cases} \tag{24}$$

We denote these interim values by v_1, v_2, v_3 and v_4 respectively.[22] The discrete grid **v** is a consequence of Bayesian updating with exponential variables. Given exponential liquidity shocks W_b and W_s, aggregate order flows are only informative about the *possibility* of informed orders with good or bad news. For example, in region IV the valuation v_4 reflects only the fact that, given the informed's plunging strategy in (22), it is impossible to see $A_s < \bar{g}$ given bad news ψ_d, but that $A_b \geq \bar{g}$ is possible given either good news ψ_u or no news ψ_n.

[22] In a symmetric market we actually have a three–point grid with $v_2 = v_3$.

Comparative statics for the updating rule (24) are intuitive. In particular, v_4 is increasing in the probability of being informed π, in the informed trader's maximum scale $\bar{g}$ and in the exponential parameter θ which puts more weight on low liquidity shocks (thereby making extreme realizations like those in region IV more informative). The reverse is true for v_1.

The last step of the equilibrium construction is to note that, given the dependence of the end-points v_1 and v_4 on the probability of being informed, we can always choose π small enough to compress the now endogenous value grid $\mathbf{v}$ to lie within the price interval $[p_1, p_2]$. Thus, under Assumption 3, our provisional assumption of order–independent prices (used to derive the informed's plunging strategy) is internally consistent.

With this characterization of the informed's optimal strategy, the existence of a Bayesian Nash equilibrium under Assumptions 1 through 3 hinges on showing that a fixed point in the liquidity trader's coefficients β and σ can be found. This is readily done using the same logic as in Proposition 4. The only conceptual difference is that informed trading exposes limit order execution to a "winner's curse" problem. Thus, even in the absence of any price effects, informed trading affects both the welfare and the trading strategies of uninformed liquidity and value traders.

The winner's curse problem is reflected in the conditional fill ratios which now depend on both the aggregate shocks W_b and W_s and on the informed trader's information. In particular, private good (bad) news tends to lower (raise) fill ratios for limit buys

$$F_{b,1} = \begin{cases} \min\{1, \frac{\sigma_m W_s}{\beta_1 W_b + \alpha_{b,1} + \beta_1 \bar{g}}\} & \text{if } (\bar{v}, \psi_u) \\ \min\{1, \frac{\sigma_m W_s}{\beta_1 W_b + \alpha_{b,1}}\} & \text{if } (\bar{v}, \psi_n) \text{ or } (\underline{v}, \psi_n) \\ \min\{1, \frac{\sigma_m W_s + \sigma_m \bar{g}}{\beta_1 W_b + \alpha_{b,1}}\} & \text{if } (\underline{v}, \psi_d) \end{cases} \tag{25}$$

and conversely raise (lower) them for limit sells

$$F_{s,2} = \begin{cases} \min\{1, \frac{\beta_m W_b + \beta_m \bar{g}}{\sigma_2 W_s + \alpha_{s,2}}\} & \text{if } (\bar{v}, \psi_u) \\ \min\{1, \frac{\beta_m W_b}{\sigma_2 W_s + \alpha_{s,2}}\} & \text{if } (\bar{v}, \psi_n) \text{ or } (\underline{v}, \psi_n) \\ \min\{1, \frac{\beta_m W_b}{\sigma_2 W_s + \alpha_{s,2} + \sigma_2 \bar{g}}\} & \text{if } (\underline{v}, \psi_d). \end{cases} \tag{26}$$

Substituting these expressions into (17) and from there into the liquidity trader's problem (16) a fixed point can be shown to exist exists in (β, σ). Thus we have

Proposition 5 *A Bayesian Nash equilibrium exists in this market under Assumptions 1 through 4.*

With existence guranteed, we again present a few numerical examples. These examples serve two purposes. First, they are useful in gauging the reasonableness of our parametric Assumption 3 on π. And second, they provide some sense of the magnitude of the winner's curse problem on the limit order book.

All of the market parameterizations in Table 2 have terminal payoffs $\bar{v} = 22$ and $\underline{v} = 19$ with probability $h(\bar{v}) = 1/2$, a "one quarter" price grid $\{p_1, p_2\} = \{20\ 3/8, 20\ 5/8\}$, a brokerage fee $T = \$.01$, preference weight $\phi = 1$ and an expected aggregate liquidity shock $1/\theta = 500$ shares. We vary the intensity of the adverse selection problem by changing both the probability of being informed π and the size of the maximal informed position $\bar{g}$. The particular values of π and $\bar{g}$ were chosen to ensure that the resulting interim valuations v are all between p_1 and p_2 (i.e., they all satisfy Assumption 3). In essence, they involve giving the informed investor a 1 or 2 percent chance of learning information worth around \$1.50 and letting her trade close to 500 or 1,000 shares using a mix of limit and market orders. Hence, the parameter restrictions implied by Assumption 3, at least in this example, do not seem unreasonable out-of-hand.

For each equilibrium we report the expected cost of a limit buy $e_{b,1}$,[23] the liquidity buyer coefficients β_1 and β_m and the total limit buys from value traders $\alpha_{b,1}$. In contrast to Section 2, the expected limit fill ratio $E(\tilde{F}_{b,1} : \psi)$ is now conditioned on whether the informed trader has good news ψ_u, no news ψ_n or bad news ψ_d. We report all three values as a measure of the magnitude of the winner's curse. Again, the corresponding values for sell orders are symmetric.

The first market in Table 2 is the benchline case with no winner's curse problem (i.e., the probability of being informed is $\pi = 0$). Not surprisingly, increasing $\bar{g}$ and/or π exacerbates the winner's curse problem in limit order execution leading to greater reliance on market orders by liquidity traders and to less aggressive positions for value traders. Moreover, these changes are not small — suggesting that the impact of the winner's curse problem on the limit order book is significant. For example, in the worst case scenario (row 5 with $\pi = .02$ and $\bar{g} = 1000$ shares) the expected limit fill ratio given bad news is three times the fill ratio given good or no news.[24] Consequently, the number of liquidity and value limit orders is approximately 20 to 25 percent

[23]Since market buys are always filled at p_2 and since the expected payoff is, by construction, at the midpoint of the bid–ask spread or 20 1/2, we again have that the expected cost for a market order is half a "tick" (now \$.125) plus the brokerage fee T.

[24]The asymmetry in expected fill ratios given good and bad news *vis-a-vis* the no news case is because bad news leads to 946.8 more market sells (from the informed trader) in the numerator of $F_{b,1}$ whereas bad news only leads to 167.1 more limit buys in the denominator.

lower.

4 Conclusion

This paper presents a single-period microstructure model with both market and limit orders. In particular, traders may choose to use any combination of these orders. A continuum of optimizing liquidity traders, value traders and possibly a strategic informed trader have symmetric access to these orders. A quadratic liquidity objective allows us to characterize equilibria in markets with either exogenous or endogenous information revelation and to solve numerically for examples. This setting also includes non-informational bid/ask spreads and brokerage costs.

Qualitative features of our analysis include (a) a random limit order book which, together with market orders, is generated by an exact linear factor model, (b) a relatively modest role for value traders (in constrast to Rock (1990)) once even moderate brokerage costs are introduced, (c) a winner's curse problem in limit order execution with informed trading and (d) a sufficient statistic for different order types for use in Bayesian updating in the presence of an informed trader.

This model offers a rich framework in which to explore the informational and transactional properties of financial markets. It can also be extended in many directions. First, we have not exhausted the ability of this framework to accomodate institutional detail. For example, it would be possible to allow market orders (rather than the specialist) to pick off exposed limit orders and to explore the impact of coarser and finer price grids on limit order placement. Second, there are also important conceptual issues still to be addressed — most notably, relaxing the restriction on the intensity of the adverse selection problem so that order flows (or more precisely the information they reveal) can also affect which limit orders are executed.

Appendix

Proof of Proposition 1: (sketch) Clearly a limit buy at the maximum possible ask (p_4 here) is dominated (in terms of execution probabilities and cost) by a market order for the same number of shares since market orders are executed at the (potentially lower) realized ask. Hence liquidity buyers will not use such buy orders.

To see why a liquidity buyer i_b avoids market and/or limit sell orders, suppose that such orders were optimal. Given that all orders are at best "break even" in terms of cost, these sell orders must be useful only in that they help a buyer to better meet his target w. Thus it must be that in some states his buy orders lead to the purchase of more than $w(i_b)$ shares. However, sell orders reduce shares bought (weakly) more in "high price" states than in "low price" states, but it is precisely in the "high price" states in which his buy orders bring in relatively fewer shares. Hence it is always possible to take a "buy" profile which includes sell orders and construct a profile without selling and with reduced buying that has both lower cost and comes closer to w. Hence, sell orders are not optimal for a liquidity buyer.

The arguments for liquidity sellers are symmetric. ∎

Proof of Proposition 2: This follows from the fact that, given a fixed $w(i)$, trader i's objective (6) is quadratic and so has first-order conditions of the form $J \cdot q \geq c \cdot w(i)$ (together with the non-negativity constraints) where J is a matrix of coefficients, q is the vector of order quantities $(b_m(i), \ldots, s_4(i))$ and c is a vector of constants. The inverse J^{-1} exists given a continuous distribution for aggregate orders. Thus the constrained critical point exists and is linear in $w(i)$, is unique and is a minimum (since the unconstrained objective is quadratic).

Note that (given common beliefs about aggregate orders and the value v), ϕ is the only investor-specific parameter on which J and c depend. Thus if all investors have identical weightings ϕ, then buyers' orders (and similarly for sellers) will be identical up to their linear dependence on individual $w(i)$'s. ∎

Proof of Proposition 3: By inspection from (2) through (5), (9) and (10), each pair of total value orders $\alpha_{b,j}$ and $\alpha_{s,j}$ appears in at most only one pair of expected costs $e_{b,j}$ and $e_{s,j}$. Thus the question is whether at each limit price p_j a pair of non-negative value orders can be found which makes the corresponding pair of expected costs non-negative. In particular, there are no "spill over" effects on limit order costs at other prices.

Given a fee $T > 0$ there clearly is no problem to insure that the expected costs for limit buys at p_1 and limit sells at p_4 are non-negative since there is no corresponding fill ratio for limit orders on the other side of the market at these prices. Thus, if in the absence of value trading expected costs would be negative, simply increase value trading in these orders to drive these costs to zero.

Consider next limit buys and sells at p_2. Note first that both expected costs $e_{b,2}$ and $e_{s,2}$ can not simultaneously be negative. This follows since for $e_{b,2}$ to be negative, we must have

$$h(v_1)(p_2 - v_1) + h(v_2)(p_2 - v_2)E(\tilde{F}_{b,2}) < 0,$$

but then we cannot also have

$$h(v_1)(v_1 - p_2)E(\tilde{F}_{s,2}) + h(v_2)(v_2 - p_2) < 0$$

as is necessary for $e_{s,2} < 0$. Thus, if one of the expected costs is negative, again increase value limit orders of that type until the order's expected cost is driven to 0. The fact that expected costs are continuous in value orders together with the fact that both expect costs cannot simultaneously be negative insures then that the expected cost for the other limit order remains non-negative.

The same logic also applies to limit orders at p_3. ∎

Proof of Proposition 4: The strategy we use is to find a fixed point in the first two non-central moments (both own and cross) of the six fill ratios $\tilde{F}_{b,j}$ and $\tilde{F}_{s,j}$.

The first step is to note that the Theorem of the Maximum together with our quadratic liquidity objective (which gurantees uniqueness of the solution and its dependence on only the first two moments of the fill ratios) implies that the optimal order coefficients (β, σ) are continuous in the first and second (and no other) fill ratio moments.

The second step is to note that the value order flows are continuous in the coefficients (β, σ). This follows from the expected costs (4) and (5), the definitions of $\tilde{F}_{b,j}$ and $\tilde{F}_{s,j}$ in (2) and (3) together with (9) and (10).

The third step is to note that the realizations $\tilde{F}_{b,j}$ and $\tilde{F}_{s,j}$ and hence their moments all are in the interval $[0, 1]$ and are continuous in the coefficients (β, σ).

The fourth step is to combine steps 1 through 3 to obtain a continuous composite mapping from a vector of conjectured fill ratio moments via the liquidity and value order flows they induce back into a new vector of moments. Since the mapping is onto with a compact domain [0,1], the Brower Fixed Point Theorem gurantees that a fixed point exists.

The final step is to use Proposition 2 to then calculate the corresponding liquidity order coefficients at the fixed point moments. ∎

Proof of Proposition 5: This result follows from repeating the same steps as the in proof of Proposition 4 after recognizing that the (again bounded) non-central cross moments between the payoff v^* and the fill ratios $F_{b,1}$ and $F_{s,2}$ must also be include in the moment fixed point. ∎

References

Admati, A., and P. Pfleiderer, 1988, A Theory of Intraday Patterns: Volume and Price Variability, *Review of Financial Studies* 1, 3-40.

Angel, J., 1992, Limit versus Market Orders, working paper, Georgetown University.

Biais, B., P. Hillion and C. Spatt, 1992, An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse, working paper, Carnegie Mellon University.

Breiman, L., 1968, *Probability*, Addison-Wesley.

Byrne, R., 1993, Discrete Prices in Securities Markets, working paper, Carnegie Mellon.

Chakravarty, S., and C. Holden, 1994, An Integrated Model of Market and Limit Orders, working paper, Indiana University.

Chung, K, 1974, *A Course in Probability Theory*, Academic Press.

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

Easley, D., and M. O'Hara, 1991, Order Form and Information in Securities Markets, *Journal of Finance*46, 905-927.

Foucault, T., 1993, Price Formation in a Dynamic Limit Order Market, working paper, HEC.

Glosten, L., 1994, Is the Electronic Open Limit Order Book Inevitable?, *Journal of Finance* 49, 1127–1161.

Glosten, L., and P. Milgrom, 1985, Bid, Ask and Transaction Prices in a Specialist Market with Heterogenously Informed Traders, *Journal of Financial Economics* 21, 123-144.

Harris, L., 1994, Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems, working paper, University of Southern California.

Harris, L., and J. Hasbrouck, 1992, Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy, working paper, University of Southern California.

Kumar, P., and D. Seppi, 1993, Information and Index Arbitrage, *Journal Of Business* 67, 481–509.

Kyle, A., 1984, A Theory of Futures Market Manipulations, in R. W. Anderson, ed.: *The Industrial Organization of Futures Markets*, (Lexington Books, Lexington, MA).

Kyle, A., 1985, Continuous Auctions and Insider Trading, *Econometrica* 53, 1315-1335.

Kyle, A.,

Parlour, C., 1994, Price Dynamics in a Limit Order Market, working paper, Queen's University.

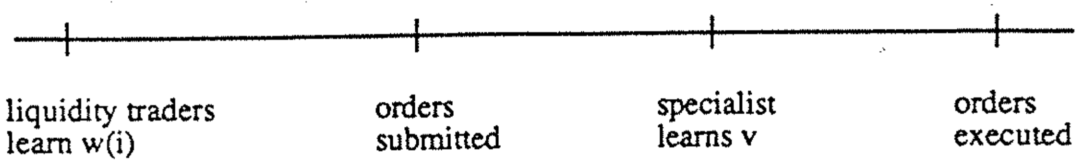

Figure 1: Timing of events

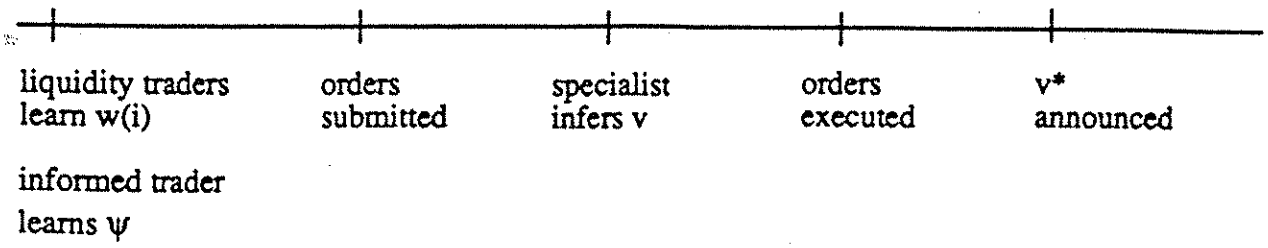

Figure 2: Timing of events

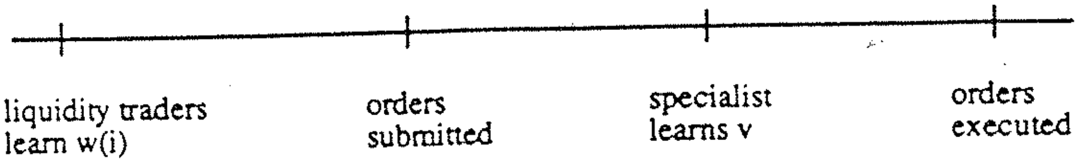

Figure 1: Timing of events

Table 1

Equilibrium with Exogenous Resolution of Uncertainty*

Probabilities			Expected Cost		Expected Fraction Executed		Liquidity Buyer's Order Coefficients			Value Traders' Buy Orders	
$h(v_1)$	$h(v_2)$	$h(v_3)$	$e_{b,1}$	$e_{b,2}$	$E(F_{b,1})$	$E(F_{b,2})$	β_1	β_2	β_m	$\alpha_{b,1}$	$\alpha_{b,2}$
Panel A: $\phi = 0.1$, Transaction Costs = \$.05 per share											
0.30	0.40	0.30	0.031	0.046	1.00	0.92	0.0000	0.1351	0.8012	0	0
0.15	0.70	0.15	0.041	0.023	1.00	0.83	0.0000	0.3043	0.6836	0	0
0.05	0.90	0.05	0.047	0.010	1.00	0.77	0.0000	0.4050	0.6171	0	0
Panel B: $\phi = 0.1$, Transaction Costs = \$.01 per share											
0.30	0.40	0.30	0.000	0.005	0.53	0.93	0.0066	0.1285	0.8664	602	0
0.15	0.70	0.15	0.001	0.000	1.00	0.44	0.0000	0.1380	0.8898	0	785
0.05	0.90	0.05	0.007	0.000	1.00	0.23	0.0000	0.1622	0.9100	0	1842
Panel C: $\phi = 0.1$, Transaction Costs = \$.001 per share											
0.30	0.40	0.30	0.000	0.000	0.05	0.79	0.0000	0.1060	0.8985	8374	167
0.15	0.70	0.15	0.000	0.000	0.11	0.24	0.0000	0.0946	0.9325	4345	1885
0.05	0.90	0.05	0.000	0.000	0.32	0.07	0.0000	0.0962	0.9505	1405	6398
Panel D: $\phi = 1.0$, Transaction Costs = \$.01 per share											
0.30	0.40	0.30	0.000	0.004	0.53	0.99	0.0007	0.0156	0.9839	687	0
0.15	0.70	0.15	0.001	0.000	1.00	0.44	0.0000	0.0153	0.9878	0	936
0.05	0.90	0.05	0.007	0.000	1.00	0.23	0.0000	0.0177	0.9902	0	2076

* All market parameterizations have exponential liquidity shocks with $\theta = \frac{1}{500}$, a value grid $\mathbf{v} = [20.0625\ 20.1875\ 20.3125]$, a price grid $\mathbf{p} = [20.000\ 20.125\ 20.250\ 20.375]$ and a convergence tolerance of 0.0001. Liquidity and value sell orders are symmetric with $\sigma_4 = \beta_1$, $\sigma_3 = \beta_2$, etc.

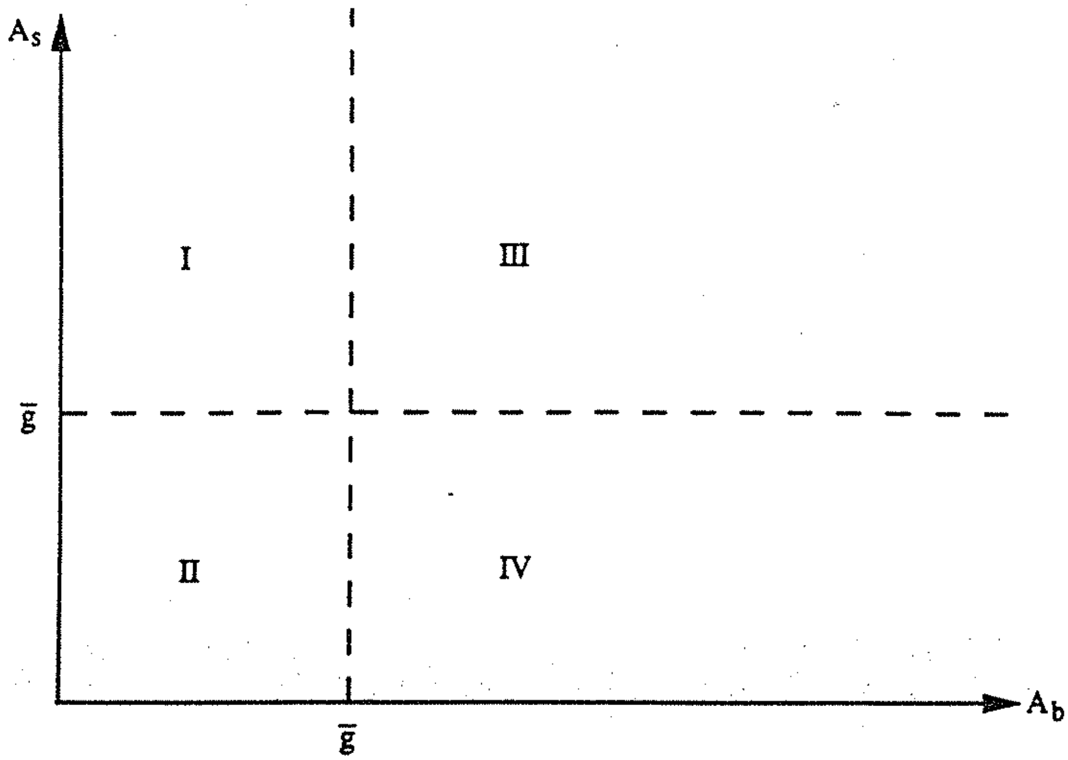

Figure 4: Valuation regions

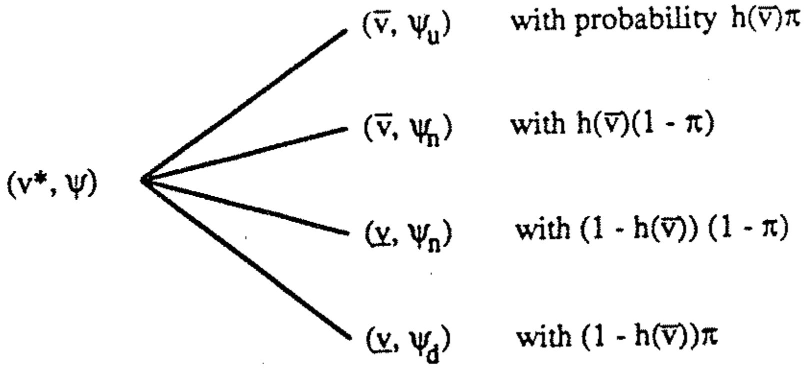

Figure 3: Possible States