Shmuel Baruch
David Eccles School of Business, University of Utah

Who Benefits from an Open Limit-Order Book?*

The NYSE opened the limit-order book to off-exchange traders during trading hours. We address the welfare implications of this change in market structure. We model a market similar to the auction that the exchange uses to open the trading day. We consider two different environments. In the first, only the specialist sees the limit-order book, while in the second the information in the book is available to all traders. We compare equilibria and find that traders who demand liquidity are better off when the book is open while liquidity suppliers are better off when the book is closed.

I. Introduction

In 2001, American security markets switched to decimal pricing. Since then, it is argued, the specialists on the New York Stock Exchange (NYSE) and the limit-order traders have been able to change quotes by offering a slightly better price (penny improvement) for a small number of shares. Thus, the inside quotes are no longer a good indicator of market conditions. Addressing the concerns of investors who desire a better look at market depth, the NYSE, as of January 24, 2002, made the limit-order book visible to the public in real time during trading hours. According the NYSE OpenBook Specification, the NYSE disseminates a full view of limit-order book beginning at 7:30 A.M., 2 hours before the market opens. In this paper, we develop a model to address the welfare implications of making the limit-order book visible prior to market opening.

The NYSE begins the trading day at 9:30 A.M. with a single-price call-type auction. "At the opening" market buy and "at the opening" market sell orders that accumulated while the exchange was closed are paired automatically by the Opening Automated Report Service (OARS). The imbalance is presented to the specialist, who then compares it

* I am grateful to Kerry Back for his guidance. I thank Hank Bessembinder, Phil Dybvig, Michel Habib, Eric Hughson, Kenneth Kavajecz, Pete Kyle, Mike Lemmon, Venkatesh Panchapagesan, Gideon Saar, Raj Singh, and an anonymous referee for their helpful comments. An earlier version of this paper was presented at the annual meeting of the American Finance Association, January 1998. Contact the author at finsb@business.utah.edu.

(*Journal of Business,* 2005, vol. 78, no. 4)

0021-9398/2005/7804-0005$10.00

with the limit orders that accumulated in his electronic book. The specialist finds a single price that will clear the market order imbalance as well as all the limit orders to buy (sell) at or below (above) the clearing price. However, unlike a typical auctioneer, the specialist can buy or sell for his own account.[1]

Other exchanges, including the Toronto Stock Exchange, the Paris Bourse, and the Frankfurt Stock Exchange, also begin the trading day with a single-price call-type auction. The conventional wisdom is that a single-price auction is a good way to establish a price that reflects broad interest. With an average of 10% of the daily dollar trading volume on the NYSE taking place at the opening, it seems that many investors prefer trading at the opening.[2] The NYSE also uses the single-price auction after a trading halt, when uncertainty is high, creating the need to establish a single price that aggregates diverse views of investors.[3]

To study the effects the change in the transparency of the limit-order book might have, we employ a stylized model of a specialist's single-price auction in two different environments: in one environment, the limit-order book is open; in the other, it is closed. In our model, liquidity traders submit market orders. These are paired automatically, and the market order imbalance is presented to the specialist. A finite number of strategic off-exchange limit-order traders submit price-contingent orders that are placed in the limit-order book. To study the efficiency of the price discovery process, our model incorporates a strategic informed trader who places a market order. The strategic specialist, after observing the market order imbalance and the limit-order book, sets the price that clears the market order imbalance, with the book taking precedence over the specialist.

Our results show that, when the market is large enough, opening the limit-order book is beneficial to market order traders, whether informed or liquidity. In fact, the price impact of market orders (reciprocal of depth) is lower on average when the book is open, so the cost of trading is lower on average with an open book, implying fewer price reversals after the opening. Moreover, we show that, on average, prices reveal more information when the book is open, implying lower post-open volatility. This result contrasts with the belief that efficiency of prices

1. See Stoll (1985) for an in-depth study of the economics of the specialist's roles in the NYSE.

2. See Madhavan and Panchapagesan (2000) who also report that, for low-volume stocks, the opening can count for as much as 25% of the total daily volume.

3. Empirical studies of single-price auction have been made by Stoll and Whaley (1990), Biais, Hillon, and Spatt (1999), and Madhavan and Panchapagesan (2000). Stoll and Whaley (1990) studied the opening on the NYSE. They found that prices tend to reverse around the opening, and they concluded that the immediacy suppliers do extract rents from the liquidity traders. Biais et al. (1999) studied the opening in the Paris Bourse, which is an open book environment where a disinterested auctioneer, a computer, sets the clearing price. They suggest that the preopening inductive prices converge to an efficient opening price.

comes at the expense of the liquidity traders (see, e.g., O'Hara 1995, p. 271). Our results are driven by the interaction between the two components of trading costs (the adverse selection and the transitory component), which is endogenous. When the book is open, the transitory component is lower, due to the increase in competition for liquidity provision. Thus, the informed trader trades more aggressively, releasing more of his private information. However, the decrease in the transitory component offsets the increase in the adverse selection component, so that overall trading costs are lower and prices are more informative in the open-book environment.

We also show that limit-order traders extract more rents when the book is closed, and numerical analyses indicate that the specialist, too, is better off in the closed-book environment. These results can be explained in the following way. In the closed-book environment, the specialist and the limit-order traders enjoy informational advantages. The specialist observes the complete structure of limit-order book, while the limit-order traders have partial knowledge of the book's structure; namely, each knows that the book contains his order. These advantages do not exist in the open-book environment. Furthermore, our results are robust with respect to the distribution of noise that we introduce into the book, as long as the market is large enough. Thus, our model demonstrates how pretrade transparency allows limit-order traders to compete more effectively with the specialist and consequently to reduce his monopoly rents.

A shortcoming of our model is a restriction we impose on the limit-order traders. Due to the difficulty of solving the limit-order traders' problem, we restrict those traders' strategies to the class of linear demand schedules. To study the practical importance of the restriction, we develop an unrestricted model that focuses solely on the limit-order traders' problem (i.e., no specialist and the informativeness of the order flow is taken as given). We used the unrestricted model to verify that the average outcomes of a restricted model are similar to those of the unrestricted one. Moreover, we show that, for small market orders, the open-book environment provides more liquidity. In contrast, for large market orders, the opposite is true. However, as in the restricted model, on average, the open-book environment is superior in terms of liquidity provision.

Limit-order traders provide better prices for large market orders in a closed-book environment because limit-order traders can condition their orders only on prices. Conditioned on extreme prices, limit-order traders have to consider two possibilities: the extreme price is due to either lack of depth in the book or a large market order. In the former case, limit-order traders extract high rents; while in the latter, they are likely trading against informed traders. Competition among limit-order traders drives their expected profit down, so that, in fact, they lose when they trade against a large market order.

Our paper is closely related to the growing literature on the limit-order book. While the current paper focuses on the limit-order book at the opening, most papers that model the book are interested in the discriminatory price auction that follows the opening. Whereas the trading protocol at the opening is a single-price auction, the continuous trading protocol is a discriminatory price auction. That is, a large market order is paired off with several limit orders, possibly at different prices.[4] Our model contributes to this literature in several ways. Ours is the only model in which strategic limit-order traders, a strategic market order trader (the informed trader), and a strategic specialist interact. In particular, this interaction allows us to study the strategy of a limit-order trader who knows that his actions alter the behavior of the specialist. Moreover, ours is the only study of trading into a closed, random depth limit-order book.

Our model is also related to the literature on transparency. Madhavan (1995) studied the effect of posttrade transparency. In a market without a posttrade disclosure requirement, dealers may be willing to provide better quotes to compete on order flow for its information content. Pagano and Roel (1996) studied different trading systems with different degrees of transparency. However, these systems are different in dimensions other than their degrees of transparency. Bloomfield and O’Hara (1999), and Flood et al. (1999) use experiments to study transparency in a pure dealer markets.

Boehmer, Saar, and Yu (2005) is an empirical study of the issues discussed in our work. They contrast market outcomes before and after the introduction of the OpenBook system and provide evidence to the main predictions made here. More specific, they find that, after the NYSE introduced the OpenBook system, the “ex ante” liquidity (as reflected by the book) and the “ex post” liquidity (taking into account price improvement by the specialist) improved. They also found some improvement in efficiency of prices following the introduction of the OpenBook. Interestingly, Boehmer et al. (2005) also compared the rate of alteration of limit-order traders before and after the introduction of OpenBook, finding that in the transparent environment the rate is higher. This is an indication that indeed the OpenBook facilitates effective competition between the limit-order traders. Each trader observes the book and adjusts his order. On the other hand, Madhavan, Porter and Weaver (1999) studied the 1990 move of the Toronto Stock Exchange into a

4. For example, Rock (1990), Seppi (1997), and Parlour and Seppi (2001) model the limit-order book in a transparent, continuous trading environment with competitive limit-order traders and a strategic specialist. Glosten (1994) is yet another variant of the Rock model; however, without the specialist. Parlour (1998) and Foucault (1999) model dynamic limit-order markets, i.e., without specialist, using the Glosten and Milgron framework. Chakravarty and Holden (1995) model a market in which the informed can submit limit orders; although, in their model, market makers ignore the information in the limit-order book.

transparent environment. They found higher spreads and higher volatility in the transparent environment. Madhavan et al. (1999) express the view that, in a transparent limit-order book environment, an informed trader can better place his market orders. This should result in higher gains for him at the expense of the limit-order traders. So, the argument goes, limit-order traders are reluctant to post their orders for fear of being picked off by an informed trader, resulting in a thinner limit-order book. However, our model shows that another argument can be made: regardless of transparency, one expects to find less-informed trading and less gains for informed traders when the book is thin. Because liquidity providers can better compete in a transparent environment, the limit-order book should be thicker in an open-book environment.

The paper is organized as follows. Section II describes the primitives of the model. Sections III and IV derive the linear restricted closed book equilibrium and the linear restricted open book equilibrium, respectively. Section V compares the equilibria. Section VI develops the unrestricted model. Finally, Section VII concludes the paper.

II. The Model

We consider a call market for a risky asset and a risk-free asset (numeraire) with the interest rate set to zero. At time 1, the risky asset pays $\tilde{v}$. We study an equilibrium in two different environments. In one environment, the limit order book is open, while in the other only the specialist observes the book. The characteristics of both environments are presented next, followed by a discussion.

Four types of participants are in our market. The first group consists of the liquidity traders. We do not model their behavior.[5] We denote their aggregate market orders by $\tilde{z}$. One strategic risk-neutral informed trader, who knows the realization of $\tilde{v}$, submits a market order, $\tilde{x}$.[6] The aggregate market order, $\tilde{y}$, which we sometimes call the *market order imbalance,* is equal to $\tilde{z} + \tilde{x}$.

Demand schedules are submitted to the specialist by $\tilde{N}$ strategic risk-neutral traders, who are called *limit-order traders.*[7] For the distinction

5. One could model the demand of liquidity trading. We have chosen not to do this for two reasons. First, linear equilibria fail to exist unless liquidity traders are risk averse. Risk aversion strengthens the bias in favor of an open-book environment. Second, we expect endogenous liquidity demand to sharpen our results. This is because now, with elastic liquidity demand, liquidity demands are greater in the environment with the lower cost of trading. Furthermore, one well-known stylized fact in finance, which finds ground in the asymmetric information literature too, is that higher volume reduces cost of trading (see Demsetz 1968).

6. In our risk-neutral environment, one can assume that $\tilde{v}$ is merely an unbiased estimator of the liquidation value.

7. A limit order is a single-step function. It sets the upper (lower) price at which a trader is willing to buy (sell) up to a specified quantity. It seems reasonable that, with decimal pricing, limit-order traders submit multiple limit orders, thus mimicking a demand schedule (see Kyle 1989).

between open and closed books to be meaningful, there must be some uncertainty about the book. In this paper, noise is introduced into the book through the number of limit-order traders, who are assumed to be drawn out of a pool of potential limit-order traders. Furthermore, conditional on the realization of $\tilde{N}$, each potential limit-order trader is equally likely to be present in the market. We denote by $f_i(\cdot)$ the ith limit-order trader's demand schedule with the interpretation that, at price p, $f_i(p)$ is the quantity the trader demands. It is convenient to denote the book's randomness by writing $f(\tilde{N}, \cdot) = \sum_{i=1}^{\tilde{N}} f_i(\cdot)$. We study equilibria in which all the limit-order traders make the same choice of a demand schedule.

The role of the specialist is to set a single price and clear the market. As a dealer, the specialist can buy and sell for his own account. However, he is subject to one important restriction. At the clearing price, the first transactions go to the book. This restriction prevents the specialist from setting an arbitrarily high (low) price and selling (buying) all the excess market orders.

Given the price, p, chosen by the specialist, the informed trader's profit is

$$(\tilde{v} - p)x, \tag{1}$$

and the profit of the ith limit-order trader is

$$(\tilde{v} - p) f_i(p). \tag{2}$$

The specialist receives the quantity

$$-\tilde{x} - \tilde{z} - \sum_{i=1}^{\tilde{N}} f_i(p),$$

so his profit is

$$(p - \tilde{v}) \left(\tilde{x} + \tilde{z} + \sum_{i=1}^{\tilde{N}} f_i(p) \right). \tag{3}$$

Our probability space has three independent random variables: $\tilde{v}$, $\tilde{z}$, and $\tilde{N}$. The liquidation value $\tilde{v}$ is normally distributed with mean $\bar{v}$ and variance σ_v^2. The aggregate liquidity order $\tilde{z}$ is normally distributed with mean zero and variance σ_z^2. The numer of limit-order traders, $\tilde{N}$, is a bounded positive integer-valued random variable. A lower bound on the support of $\tilde{N}$ is needed for certain results. We impose no other distributional assumptions on $\tilde{N}$. Given a random variable $\tilde{u}$, the notations u and $\bar{u}$ are used to denote its realization and its expected value, respectively.

Due to the mathematical difficulty of solving the limit-order traders' problem, we can use only approximation methods. Our approach is to solve analytically an approximate model: we restrict the limit-order traders to linear demand schedules.[8]

Our model does not incorporate the group of floor brokers because we want a level playing field. The floor traders provide an advantage for their clients whether the book is closed or open. In a closed-book environment, they can communicate information from the floor to their off-exchange clients, in particular, tell them what is in the limit-order book. On the other hand, in an open limit-order book environment they can "work" the orders of their clients rather than posting limit orders. Interestingly, Sofianos and Werner (1997) found that the participation of floor brokers at the opening is very low. They estimated the value of floor brokers' executed orders at the opening, excluding orders submitted through the OARS, to be only 0.9%.[9]

III. Closed Book

In the closed-book environment, the informed trader can condition his market order only on the asset value $\tilde{v}$. He has no information about the book when he submits his order. We therefore write his market order as a function $x(v)$.

An important feature of a closed-book environment is that a limit-order trader knows his own demand schedule, f_i, and thus possesses some information on the book's content. Here, this information is captured by the fact that a limit-order trader knows that he is active in the market; that is, the trader knows that the book contains his order. Let m_i be the indicator function of the event that the ith trader is active; that is, $m_i = 1$ when he is active and $m_i = 0$ otherwise.

The specialist observes both the market order imbalance

$$\tilde{y} \equiv x(\tilde{v}) + \tilde{z}$$

and the book $f(\tilde{N}, \cdot)$ before choosing the price p, so he chooses the price as a function of $\tilde{y}$ and $f(\tilde{N}, \cdot)$. We write this function as $P(y, f)$.

8. Using the unrestricted model presented in Section VI, we can show that the smaller the information content in order flow, the smaller the expected price deviation a limit-order trader expects and, hence, the better the linear approximation is. This result is available from the author on request. We also show (see corollary 3) that, when the number of limit-order traders is very large, our equilibrium outcomes approach the outcomes of the competitive and unrestricted model, and we obtain this convergence result even for large expected price deviations.

9. Also Boehmer et al. (2005) found that, after OpenBook was introduced, volume attributed to floor brokers declined relative to volume attributed to the limit-order book.

The informed trader's expected profit from a market order x, contingent on a realization v of the random variable $\tilde{v}$, is

$$E\{v - P[x + \tilde{z}, f(\tilde{N}, \cdot)]x\}. \tag{4}$$

The ith limit-order trader takes the demand schedules of the other limit-order traders as given. It is convenient to focus on the decision problem of the first trader, since all the limit-order traders face the same decision problem. Given f_i for $j > 1$, set

$$f_{-1}(\tilde{N}, p) = \sum_{j=2}^{\tilde{N}} f_j(p).$$

Given a demand schedule f_1, the book is

$$f_1(\cdot) + f_{-1}(\tilde{N}, \cdot).$$

The expected profit of the limit-order trader, contingent on the knowledge that he is active in the market, is

$$E[(\tilde{v} - \tilde{p}) f_1(\tilde{p}) | m_i = 1], \quad \text{where } \tilde{p} \equiv P[\tilde{y}, f_1(\cdot) + f_{-1}(\tilde{N}, \cdot)] \tag{5}$$

Given a market order y and a book $f(\tilde{N}, \cdot)$, the specialist chooses the price p to maximize

$$E\{(p - \tilde{v})[y + f(\tilde{N}, p)] | f(\tilde{N}, \cdot), x(\tilde{v}) + \tilde{z} = y\}. \tag{6}$$

A linear restricted equilibrium consists of a decision rule $x(v)$ for the informed trader, a demand schedule f_1 for each of the limit-order traders, and a decision rule $P(y, f)$ for the specialist such that

1. The market order $x(v)$ maximizes (4) for each realization v of $\tilde{v}$.
2. The demand schedule f_1 maximizes (5) over the class of linear functions.
3. The price rule $P(y, f)$ maximizes (6) for each realization y of $x(\tilde{v}) + \tilde{z}$ and each linear demand schedule f.

Theorem 1. If $E\left[\frac{\tilde{N}-2}{\tilde{N}^3} \middle| m_i = 1\right] > 0$, then there exists a linear restricted equilibrium (hereafter, equilibrium) in which[10]

(i) The ith limit-order trader's demand schedule has the form

$$f_i(p) = (\tilde{v} - p)B_c.$$

10. We use the subscript c to indicate the closed-book environment.

(ii) The informed trader's decision rule has the form

$$x(\tilde{v}) = \beta_c(v - \bar{v}),$$

(iii) The specialist's price rule has the form

$$P(y, f) = \bar{v} + \frac{1}{2}\left[b_c + \frac{1}{-f'(p)}\right]y + \frac{1}{2}\frac{f(\bar{v})}{-f'(p)}$$

In particular, since $f'(p) = NB_c$ and $f(\bar{v}) = 0$, the pricing rule in equilibrium simplifies to $p = B - V + \tilde{\lambda}_c y$ where $\tilde{\lambda}_c \equiv \frac{1}{2}\left(b_c + 1/\tilde{N}B_c\right)$.

The triple (B_c, b_c, β_c) is given as the positive solution of the following system of equations:

$$\begin{cases} b_c = \frac{\beta_c \sigma_v^2}{\beta_c^2 \sigma_v^2 + \sigma_z^2} \\ \tilde{\lambda}_c = \frac{1}{2}\left(b_c + \frac{1}{B_c}\frac{1}{\tilde{N}}\right) \\ B_c = \frac{1}{b_c}\sqrt{E\left[\frac{\tilde{N}-2}{\tilde{N}^3}\middle| m_i = 1\right]} \\ \beta_c = \frac{1}{2E\tilde{\lambda}_c}. \end{cases} \qquad (7)$$

Proof. Before we show that the system (7) defines an equilibrium, we need to show that the system possesses a solution. If a solution exists, then it implies that

$$\begin{cases} b_c = \frac{\beta_c \sigma_v^2}{\beta_c^2 \sigma_v^2 + \sigma_z^2} \\ E\tilde{\lambda}_c = \frac{1}{2}\left(b_c + \frac{1}{B_c}E\frac{1}{N}\right) \\ B_c = \frac{1}{b_c}\sqrt{E\left[\frac{\tilde{N}-2}{\tilde{N}^3}\middle| m_i = 1\right]} \\ \beta_c = \frac{1}{2E\tilde{\lambda}_c}. \end{cases} \qquad (8)$$

It is straightforward to see that a unique positive solution to this system exists. Endowed with $E\tilde{\lambda}_c$, we can solve the system (7), where the primitives are σ_v, σ_z, $E\left[\frac{\tilde{N}-2}{\tilde{N}^3}\middle| m_i = 1\right]$, $E\tilde{\lambda}_c$, and the realization N of $\tilde{N}$.

The proof that the system (7) defines a linear restricted equilibrium is given in Appendix A. Q.E.D.

The assumption that $\tilde{N} \geq 2$ and is nondegenerate is sufficient for the existence of an equilibrium. It is, however, not necessary. What is important is that a limit-order trader does not assign too much weight to the event that he has monopoly power, that is, the event $\{\tilde{N} = 1\}$.

The equilibrium we found has several features that distinguish it from what has been done so far in the literature. Here, not only does a strategic limit-order trader utilize information in the clearing price by conditioning his demand on the opening price, he also takes into account the strategy of the specialist who chooses his position only after all the orders have been submitted to him.[11] Furthermore, because the number of traders in our model is uncertain, the price impact of a market order, measured by $\tilde{\lambda}_c$, is random. Neither a limit-order trader nor the informed trader observes $\tilde{\lambda}_c$, although, as we mentioned, a limit-order trader possesses some information about it. It turns out, as the following lemma demonstrates, that there is a simple way to express the statistical value of that information which we denote by m_i.

LEMMA 1. The ratio of conditional to unconditional probabilities of $\tilde{N}$ is

$$\frac{\text{Prob}(\tilde{N} = N | m_i = 1)}{\text{Prob}(\tilde{N} = N)} = \frac{N}{E\tilde{N}}.$$

In particular, for any $g(\cdot)$,

$$E\left[g(\tilde{N}) | m_i = 1\right] = \frac{Eg(\tilde{N})\tilde{N}}{E\tilde{N}}.$$

Proof. Since each potential limit-order trader is chosen with the same probability out of the pool of potential limit-order traders, we have $\text{Prob}(m_i = 1 | \tilde{N} = N) = N/K$, where K is the number of potential traders.[12] This implies that

$$\begin{aligned}\text{Prob}(m_i = 1) &= \sum_N \text{Prob}(m_i = 1, \tilde{N} = N) \\ &= \sum_N \text{Prob}(m_i = 1 | \tilde{N} = N)\text{Prob}(\tilde{N} = N) \\ &= \frac{E\tilde{N}}{K}\end{aligned}$$

and hence

$$\begin{aligned}\text{Prob}(\tilde{N} = N | m_i = 1) &= \text{Prob}(\tilde{N} = N)\frac{\text{Prob}(m_i = 1 | \tilde{N} = N)}{\text{Prob}(m_i = 1)} \\ &= \text{Prob}(\tilde{N} = N)\frac{N}{E\tilde{N}}.\end{aligned}$$

Q.E.D.

11. Other models that model a strategic specialist, such as Rock (1990), assume the limit-order traders are nonstrategic.
12. K is the upper bound on the support of the distribution of $\tilde{N}$, which we assume to exist.

Intuitively, we expect that the larger is the number of limit-order traders, the less valuable the information a limit-order trader has. Indeed, the lemma shows that the larger the values $\tilde{N}$ can take, the closer to 1 is the ratio of conditional to unconditional probabilities. However, the equilibrium outcomes are determined by aggregation. Thus, even with a large expected number of limit-order traders, we cannot rule out the informational advantage limit-order traders possess in a closed-book environment. The conditional expectation that appears in (8), $E[(\tilde{N}-2)/\tilde{N}^3|m_i=1]$, is equal to $E(\tilde{N}-2)/\tilde{N}^2 1/EN$. It is convenient to rewrite the system of equations (8) as

$$\begin{cases} b_c &= \frac{\beta_c\sigma_v^2}{\beta_c^2\sigma_v^2+\sigma_z^2} \\ E\lambda_c &= \frac{1}{2}\left(b_c+\frac{1}{B_c}E\frac{1}{\tilde{N}}\right) \\ B_c &= \frac{1}{b_c}\sqrt{E\frac{\tilde{N}-2}{\tilde{N}^2}\frac{1}{E\tilde{N}}} \\ \beta_c &= \frac{1}{2E\lambda_c}. \end{cases} \tag{9}$$

Lemma 1 helps us gain some insight into the equilibrium in the closed-book environment. The lemma implies that, whenever N is greater than $E\tilde{N}$, the conditional probability of $\tilde{N}$ with respect to the event $\{m_i=1\}$ assigns more weight to the event $\{\tilde{N}=N\}$ than does the unconditional probability. It follows that each of the limit-order traders expects the price impact of a market order to be smaller than its unconditional average. Indeed, from the second equation in (9),

$$E[\tilde{\lambda}_c\,|\,m_i=1]=\frac{1}{2}\left[b_c+\frac{1}{B_c}E\left(\frac{1}{N}\middle|m_i=1\right)\right]=\frac{1}{2}\left(b_c+\frac{1}{B_c}\frac{1}{EN}\right)\le E\tilde{\lambda}_c,$$

where the second equality follows from lemma 1 with $g(N)=1/N$.

The liquidity that a limit-order trader provides is inversely related to his belief about the aggregate liquidity provided by the market. One could argue that, since limit-order traders overestimate aggregate liquidity (i.e., underestimate $\tilde{\lambda}_c$), opening the book should increase liquidity provision. To make this statement precise, we need first to know how the specialist and the informed trader will revise their strategies in response to opening the book. This is the aim of the next section.

IV. Open Book

In this section, we would like to remove some of the specialist's informational advantage by opening the book. To be consistent with the NYSE OpenBook specifications, the specialist does not disclose the market-order imbalance. According the the NYSE, "In some cases, market orders comprise the majority of pre-opening interest, and market order imbalances become the key determinant to where a stock will open."[13] Thus, the book alone cannot indicate the opening price.[14]

Modeling the dynamic of an open-book environment is a complicated task. Instead, the approach taken in this paper is to assume that when the market is called the book is in a state of equilibrium; that is, given the book's status, no single limit-order trader desires to change his order. We continue, as in the closed-book environment, to maintain the role of the specialist as the "follower," who takes his actions only after the book has reached equilibrium. This time, however, the specialist has no informational advantage, since everyone sees the book before the market is called.

A book in a state of equilibrium is the one that results from a static Bayesian Nash equilibrium in pure strategies under the assumption that N is common knowledge. In such an equilibrium, each of the traders perfectly predicts the book's structure before submitting his order. In particular, once the book is realized, no trader desires to change his order, and the specialist can call the market, that is, announce the price and clear the market.

Therefore, we consider an equilibrium where the informed trader's market order is a function $x(v, N)$, a demand schedule is a linear function $f_1(N, \cdot)$, and the price rule is $P(N, y, f)$. The informed trader's expected profit from a market order x is

$$E\{v - P[N, x + \tilde{z}, f(N, \cdot)]x\}. \tag{10}$$

The expected profit of the limit-order trader is

$$E(\tilde{v} - \tilde{p}) f_1(\tilde{p}), \quad \text{where } \tilde{p} \equiv P[N, \tilde{y}, f_1(\cdot) + f_{-1}(N, \cdot)]. \tag{11}$$

Given a market order y and a book f, the specialist expected profit is

$$E\{(p - \tilde{v})[y + f(p)] | x(\tilde{v}, N) + \tilde{z} = y\}. \tag{12}$$

A linear restricted equilibrium consists of a decision rule $x(v, N)$ for the informed trader, a decision rule $f_1(N, \cdot)$ for each of the limit-order traders, and a decision rule $P(N, y, f)$ for the specialist such that

1. The market order $x(v, N)$ maximizes (10) for each realization v of $\tilde{v}$.
2. The demand schedule $f_1(N, \cdot)$ maximizes (11) over the class of linear functions.
3. The price rule $P(N, y, f)$ maximizes (12) for each realization y of $x(N, \tilde{v}) + \tilde{z}$ and each linear demand schedule f.

13. See www.nysedata.com/openbook/FAQ.htm.

14. This is in contrast with the Paris Bourse, where each time a new order is placed, a new inductive price is announced.

THEOREM 2. If $N > 2$, then there exists a linear restricted equilibrium (hereafter, equilibrium) in which[15]

(i) The demand schedule is given by

$$f_1(N, p) = (\bar{v} - p)B_o(N).$$

(ii) The informed-trader decision rule is given by

$$x(\tilde{v}, f) = \beta_o(N)(\tilde{v} - \bar{v}).$$

(iii) The price rule has the form

$$P(N, y, f) = \bar{v} + \frac{1}{2}\left[b_o(N) + \frac{1}{-f'(p)}\right]y + \frac{1}{2}\frac{f(\bar{v})}{-f'(p)}.$$

In particular, in equilibrium, $p = \bar{v} + \lambda_o(N)y$, where

$$\lambda_o(N) \equiv \frac{1}{2}\left[b_o(N) + \frac{1}{NB_o(N)}\right].$$

The triple $(B_o(N), \beta_o(N), b_o(N))$ is the positive solution of the following system of equations:

$$\begin{cases} b_o = \frac{\beta_c \sigma_v^2}{\beta_o^2 \sigma_v^2 + \sigma_z^2} \\ \lambda_o = \frac{1}{2}\left(b_o + \frac{1}{NB_o}\right) \\ B_o = \frac{1}{b_o}\sqrt{\frac{N-2}{N^3}} \\ \beta_o = \frac{1}{2\lambda_o}. \end{cases} \tag{13}$$

Proof. This is merely a special case of theorem 1, in which the distribution that governs $\tilde{N}$ is degenerate. Indeed, once we consider $\tilde{N}$ as known, system (7) reduces to (13). Q.E.D.

Despite the pretrade transparency, the semi-strong-efficient condition, $\tilde{p} = E[\tilde{v} \mid \tilde{p}]$, does not hold in equilibrium because of the market power of the liquidity providers. In fact, the specialist's and the value traders' expected gains are strictly positive. However, we can prove the following.

COROLLARY 3. In the limit, as the lower bound of N goes to infinity, the equilibrium in the open-book environment converges to the one found in Kyle (1985). In particular, in the limit the specialist acts as an auctioneer.

Proof. The market efficiency condition holds if and only if $\lambda_o = b_o$ (see Kyle 1985). From the second equation in (13), this condition holds

15. We use the subscript o to indicate the open-book environment.

if the competition among the value traders results in $1/NB_o = b_o$. It follows from the specialist price rule that, in that case, the specialist takes no position. From the third equation in (13), it follows that

$$\frac{1}{NB_o} = \sqrt{\frac{N}{N-2}}\, b_o > b_o.$$

However, as N goes to infinity, $\sqrt{N/(N-2)}$ goes to 1 and prices become efficient. Q.E.D.

V. Comparison of Equilibria

Due to the risk-neutrality assumption, the model we presented is a zero-sum game. Hence, moving from one environment to the other cannot benefit everyone. In this section, we determine who gains from the closed-book environment and who gains from moving to the open-book environment.

It is convenient to introduce the change of variables,

$$\begin{aligned} \tilde{a} &:= \frac{1}{\tilde{N}} \\ \tilde{r} &:= \tilde{a} - 2\tilde{a}^2 = \frac{\tilde{N}-2}{\tilde{N}^2}, \end{aligned} \tag{14}$$

and express the solution of the closed book equilibrium (system [9]) in terms of $E\tilde{N}$, $E\tilde{a}$, and $E\tilde{r}$:[16]

$$\begin{cases} b_c &= b(E\tilde{N}, E\tilde{a}, E\tilde{r}) \\ \beta_c &= \beta(E\tilde{N}, E\tilde{a}, E\tilde{r}) \\ E\tilde{\lambda}_c &= \lambda(E\tilde{N}, E\tilde{a}, E\tilde{r}) \\ B_c &= B(E\tilde{N}, E\tilde{a}, E\tilde{r}). \end{cases}$$

We note that the same functional form of the right-hand side can be used to express the realization of the open book equilibrium, that is, system (13):

$$\begin{cases} \tilde{b}_o = b(\tilde{N}, \tilde{a}, \tilde{r}) \\ \tilde{\beta}_o = \beta(\tilde{N}, \tilde{a}, \tilde{r}) \\ \tilde{\lambda}_o = \lambda(\tilde{N}, \tilde{a}, \tilde{r}) \\ \tilde{B}_o = B(\tilde{N}, \tilde{a}, \tilde{r}). \end{cases}$$

16. There is no ambiguity regarding σ_v and σ_z. Hence, we treat them as parameters and omit them.