

Price Discovery in Auction Markets: A Look Inside the Black Box

Ananth Madhavan

University of Southern California

Venkatesh Panchapagesan

Washington University in Saint Louis

Opening mechanisms play a crucial role in information aggregation following the overnight nontrading period. This article examines the process of price discovery at the New York Stock Exchange single-price opening auction. We develop a theoretical model to explain the determinants of the opening price and test the model using order-level data. We show that the presence of designated dealers facilitates price discovery relative to a fully automated call auction market. This is consistent with specialists extracting information from observing the evolution of the limit order book. In addition, the specialist's opening trade reflects noninformational factors such as price stabilization requirements.

In his classic description of trading on the Amsterdam Stock Exchange, Joseph de la Vega¹ wrote

On geographical maps fine dotted lines are drawn around undiscovered regions which are named *Terra Incognita*. On the Exchange, too, there are many secret operations which I have not been able to discover. . . .

While our knowledge of geography has advanced dramatically over the three centuries since De La Vega's comparison, many fundamental issues concerning the operation of financial markets are not well understood. In particular, the process by which securities markets translate investors' latent demands into realized prices and volumes remains unclear. This article analyzes the workings of this "black box" using order-level data from the New York Stock Exchange (NYSE) single-price opening auction.

We thank Yakov Amihud, Bruno Biais, Ian Domowitz, Tom George, Larry Harris, Joel Hasbrouck, Frank Hatheway, Maureen O'Hara, Patrik Sandas, George Sofianos, Chester Spatt, and an anonymous referee for their helpful comments. Seminar participants at Boston College, University of California at Irvine, University of Michigan, Rice University, Stanford University, University of Texas at Austin, University of Washington, NBER Market Microstructure Conference, JFI Symposium on the Design of Financial Systems and Markets, Western Finance Association Meetings, CEPR European Summer Symposium on Financial Markets, European Finance Association meetings, and the NYSE-SBF Bourse de Paris Conference on Global Equity Markets provided many useful suggestions. We also particularly thank Larry Glosten (the editor) for several constructive suggestions that have improved the quality of this article. The comments and opinions contained in this article are those of the authors alone. Of course, any errors are entirely our own. Address correspondence and reprint requests to *Venkatesh Panchapagesan*, Olin School of Business, Washington University in St. Louis, One Brookings Dr., St. Louis, MO 63130, or e-mail: panchapagesan@olin.wustl.edu.

¹ *Confusion de Confusiones: Portions Descriptive of the Amsterdam Stock Exchange* (1688) by Joseph de la Vega. Translation by H. Kellenbenz, Baker Library, Harvard University (1957).

This issue is important for academic and practical reasons. Trading mechanisms vary in many dimensions, including in their reliance on dealers to provide liquidity and the extent to which they provide information on the trading process to investors. The impact of these differences on price efficiency and liquidity is a topic of considerable academic importance, and one that we address in this article. Interest in the NYSE opening is also motivated by the widespread, albeit controversial, belief that such multilateral trading systems are efficient mechanisms to aggregate diverse information.² Indeed, many *continuous* markets use single-price auction mechanisms at the open, close, or to reopen following a trading halt.³ Consequently, there is interest in how call auctions operate and whether such systems can be used more widely to trade securities.

The NYSE's single-price opening auction is also of considerable practical interest in several dimensions. First, the opening is an especially crucial period because uncertainty regarding fundamental values is high following the overnight or weekend nontrading period. Second, unlike many automated systems, a designated dealer or specialist determines the opening price. The U.S. Securities and Exchange Commission (1963) notes that "the control that specialists have on prices is nowhere better illustrated than at openings." The impact of specialist intervention at the open is highly relevant to other markets considering the adoption of similar protocols at the open, but has not been examined in detail. Finally, there is growing concern regarding delayed or difficult openings in other equity markets. Wyatt (1999) notes, "Some big traders and institutional investors say the way Nasdaq opens trading each day has bred chaos, and they are calling for change." The NASD is considering a single-price opening as a way to resolve the problems of locked and crossed markets [see Cao, Ghysels, and Hatheway (2000)] that often characterize the opening. Our article sheds light on the potential impacts of such initiatives.

We develop a model of a single-price call auction where a strategic market maker (or specialist) sets the opening price after observing the limit order book. The model closely resembles opening procedures on the NYSE and the Frankfurt Stock Exchange, among others.⁴ The ability to observe the evolution of the limit order book conveys valuable information to the specialist. Consequently, if the specialist were to trade on this information strategically, the opening price will be more efficient than the price that would otherwise prevail in an auction without dealer intervention. However, specialist trading

² Madhavan (1992) provides a theoretical argument for batch markets as a way to reduce market failures caused by information asymmetries; Economides and Schwartz (1995) and Schwartz and Wood (1995) argue for wider use of call auction markets, for example, to facilitate trading on derivative expiration dates.

³ Meier (1998) reports that at year-end 1997, 71% of exchanges used special procedures at the open and 51% use batch auctions following a trading halt. Recently, the Paris Bourse introduced a closing call auction to alleviate problems induced by large volumes at the close.

⁴ The Frankfurt Stock Exchange opens with a call auction conducted by a designated dealer or Makler [see, e.g., Kehr, Krahenen, and Theissen (1998)], much like the NYSE.

at the open is also affected by noninformation motives such as inventory control or exchange obligations, factors that might add noise to the opening price. Thus the net effect of dealer trading on price efficiency is an empirical question.

We then directly test the model's predictions regarding the specialist's strategic trading using data on *orders* submitted for the open for a wide cross section of NYSE stocks. Order-level data are especially important for our analyses because the opening price is affected by all *latent* demands, not just orders that ultimately result in transactions. We use these data to compute the market-clearing price that would have prevailed in the absence of trading by floor participants. We compare this statistic to the actual opening price to make inferences about the information content of the limit order book and the ability of dealers to affect market prices.

The empirical analysis yields several interesting results. We find support for our model of dealer behavior at the open. In particular, the specialist's opening actions reflect expectations of future price movements and price stabilization. Surprisingly, given the specialist's advantages at the open, there is no evidence that the specialist uses this time to rebalance inventory. We provide evidence from regression estimates and variance ratio tests that the specialist sets a more efficient price than the price that would prevail in a pure call market using only public orders. Finally, a market clearing price based solely on public orders need not always exist. This notional price also exhibits considerable variation in less actively traded stocks; relatively small order imbalances can generate prices that are not economically meaningful. It may not be practical to trade all securities—especially ones that are thinly traded—in a fully automated call market.

This article contributes to the literature in several ways. Our theoretical analysis provides insight into the operation of call markets [see Ho, Schwartz, and Whitcomb (1985) and Wohl and Kandel (1997)] and the role of specialists [see Madhavan and Smidt (1993)] in determining prices. It also sheds light on the relative merits of floor versus automated trading systems and the effects of transparency, as discussed by Benveniste, Marcus, and Wilhelm (1992) and Domowitz and Wang (1994). Our empirical analysis is related to work by Garbade and Sekaran (1981), Brooks and Su (1997), Cao, Ghysels, and Hatheway (2000), Kehr, Krahnen, and Theissen (1998), and Biais, Hillion, and Spatt (1999) on exchange openings. It also complements analyses by Amihud and Mendelson (1987, 1991), Stoll and Whaley (1990), Forster and George (1996), and George and Hwang (1997). Finally, these results provide additional empirical support for experimental studies [Bloomfield and O'Hara (1999) and Schnitzlein (1996)] where mechanism design, and transparency in particular, significantly affects the process of price formation.

We proceed as follows. Section 1 develops a theoretical model that serves as the basis for our subsequent empirical analysis. Section 2 describes the

institutional framework and our data sources and procedures. Section 3 contains our empirical analyses of the opening and the effects of dealer intermediation. Section 4 concludes.

1. The Analytical Framework

1.1 Notation and assumptions

In this section we develop a theoretical model of a single-price auction that yields explicit empirical hypotheses. We consider a market in which a single risky asset with unknown value is traded in a call auction mechanism. Trading is modeled as a two-stage game. In the first stage, public investors submit price-contingent orders for execution at the open.⁵ In the second stage, a designated dealer (or specialist) views the entire order book and strategically selects a single opening price, accommodating any excess demand from his or her own inventory. To focus attention on the nature of opening protocols, we do not model the subsequent continuous market.

Public investors are assumed to have negative exponential expected utility functions of the form $u(W_i) = -e^{-\rho_i W_i}$, where, for investor i , $\rho_i > 0$ is the coefficient of risk aversion and W_i is the investor's terminal wealth. Denote by v the stock's fundamental or liquidation value, p the opening price, e_i the investor's initial share endowment (with negative values indicating short positions), q_i the number of shares purchased (with positive values indicating purchases and negative values indicating sales), and c_i the investor's initial cash position. Then,

$$W_i = v(q_i + e_i) + c_i - pq_i, \quad (1)$$

where we have normalized the interest rate on the riskless asset to zero.

There are two types of public investors: informed traders and uninformed (liquidity) traders. Informed traders obtain a private information signal about the liquidation value of the asset. Suppose the prior distribution of the unknown asset value v is normal with mean μ and precision (the inverse of the variance) ζ . Informed traders receive a signal s drawn from a normal distribution with mean v and precision ψ .⁶ Let Ω_i denote the information set of informed trader $i = 1, \dots, N$; using the properties of the normal distribution, conditional upon Ω_i , trader i views v as distributed normally with mean $v_0 = E[v|\Omega_i] = \mu\gamma + s(1-\gamma)$, where $\gamma = \zeta/(\zeta + \psi)$, and conditional variance $\sigma^2 = \text{var}[v|\Omega_i] = (\zeta + \psi)^{-1}$. Informed traders are assumed to be

⁵ Prices and quantities are assumed to be continuous so investors submit price functions. See, for example, Glosten (1994).

⁶ We can extend the model to incorporate heterogeneous information without altering our qualitative results.

price-takers, an assumption we discuss in more detail below. Maximizing expected utility is equivalent to maximizing the certainty equivalent

$$v_0(q_i + e_i) + c_i - pq_i - \left(\frac{1}{2}\right)\rho_i\sigma^2(q_i + e_i)^2. \quad (2)$$

Thus an informed trader submits an opening order that is a linear function of price

$$q_i(p) = a_i - b_i p, \quad (3)$$

where $a_i = (v_0)/(\rho_i\sigma^2) - e$ and $b_i = 1/(\rho_i\sigma^2)$. Observe that the intercept reflects not only the trader's fundamental value but also the endowment hedging motive, while the slope of the demand function is inversely related to the trader's risk aversion and to the uncertainty regarding private information. Thus the order does not fully reveal the trader's information signal.

In addition to the N informed traders, $K \geq 0$ uninformed investors trade for noninformation or liquidity-based reasons, such as life-cycle consumption needs. We denote the trade of an uninformed trader j by x_j with the sign convention that positive values imply share purchases while negative values imply sales. Consistent with the previous literature, we assume the demands of liquidity traders are exogenous, so that they trade using market orders. In what follows, we assume that traders' initial endowments are drawn from a normal distribution with a constant mean, normalized to zero. Note, however, that K may be zero in our model, because the hedging demands of informed traders provides the noise necessary to avoid market failure.

1.2 An automated call market

Before turning to the second stage of the game where a strategic specialist sets the price, it is useful to examine an automated or fully electronic call market [see, e.g., Domowitz and Wang (1994)] without a designated dealer. Traders submit their orders electronically to an automated system that determines the price at which aggregate excess demand is zero, that is, the classical *Walrasian* price.

From Equation (3), the aggregate excess demand from public orders can be written as a function of price

$$Q(p) = \sum_i^N q_i(p) + \sum_j^K x_j \equiv (v_0 - p) \sum_i^N b_i - \sum_i^N e_i + \sum_j^K x_j. \quad (4)$$

Let p^* denote the market-clearing price in the automated call market, defined by the equation $Q(p^*) = 0$. From Equation, we obtain

$$p^* = v_0 + \frac{(\sum_j^K x_j - \sum_i^N e_i)}{\sum_i^N b_i} \equiv v_0 + \omega. \quad (5)$$

The market-clearing price is equal to value plus a mean-zero noise term, ω , that captures the effects of hedging of endowment risk and liquidity trading. It follows that the market-clearing price p^* is an unbiased estimator of true value.

1.3 A call market with a designated market maker

1.3.1 Market-maker behavior. In some opening mechanisms, including that used by the NYSE, there is an additional stage of the game where a designated dealer or specialist selects the opening price (which need not necessarily be p^*) and absorbs any excess demand or supply from inventory. To describe the process by which the specialist selects the price, we need to define more precisely the specialist's objective function. We assume that there are no additional participants in the two-stage game. We can relax this assumption to allow participation by floor traders, a point that we discuss later.

Like public investors, the specialist maximizes a negative exponential expected utility function. However, the specialist's decision problem differs from a public investor in important respects. First, the specialist is not a price taker instead viewing price as a choice variable. Second, the specialist is expected to "establish a fair opening price close to the prior day's last sale" to maintain price continuity [U.S. Securities and Exchange Commission (1963)]. There is no explicit rule about what constitutes continuity, but it is one of several criteria by which specialists are periodically evaluated. Consequently, substantial deviations between opening and closing prices are likely to adversely affect the specialist's reputational capital. Third, unlike public investors, the specialist sees the individual orders that constitute the aggregate excess demand function. As we will see, the ability to observe the evolution of the limit order book confers an informational advantage over public investors.

Let W_s denote the specialist's terminal wealth and let z denote the specialist's trade, where $z > 0$ denotes specialist purchases and $z < 0$ denotes specialist sales. We assume that the specialist faces costs from failing to maintain price continuity of the form $-\delta(p_0 - p^c)^2$, where $\delta > 0$ is a constant and p^c denotes the previous day's closing price. These costs are understood to represent the implicit costs arising from, say, the loss of reputational capital and the consequent loss of profits should the stock be reallocated by the exchange to another specialist. Then the specialist's post-trade wealth is given by

$$W_s = v(z + e_s) + c_s - p_0 z - \delta(p_0 - p^c)^2, \quad (6)$$

where c_s represents holdings of the riskless asset, e_s is the specialist's share inventory prior to the open, and $p_0(z)$ represents the opening price on the specialist's trade.

1.3.2 Specialist's beliefs and the limit order book. To describe the optimal actions of the specialist, we first need to describe the information content of the limit order book. Recall that the prior distribution of v is normal with mean μ and precision (the inverse of the variance) ζ . Suppose that observing the evolution of the limit order book provides the specialist with an information signal about asset value. Suppose further that this signal is normally distributed with mean v and precision π_s , whose realization is $\bar{y}$. Let Ω_s denote the information set of the specialist given the limit book. Using the properties of the normal distribution, the specialist's posterior distribution of v is also normal with mean $v_s = E[v|\Omega_s] = \mu\chi + \bar{y}(1 - \chi)$, where $\chi = \zeta/(\zeta + \pi_s)$. The conditional variance of the specialist's posterior is denoted by $\theta^2 = \text{var}[v|\Omega_s] = (\zeta + \pi_s)^{-1}$.

The specialist's signal is derived from observing the evolution of the limit order book, that is, seeing the individual orders as they arrive. Suppose initially that $K = 0$, that is, that there are no liquidity traders. An individual demand function provides a noisy signal of the informed trader's signal because the portfolio hedging component of trade cannot be distinguished from the information component. Upon observing order i , the specialist can form a statistic y_i where.

$$y_i = \frac{a_i b_i^{-1} - \mu\gamma}{1 - \gamma} = s - \frac{e_i}{b_i(1 - \gamma)}. \quad (7)$$

This statistic provides a noisy signal about the informed trader's private signal about the asset's value, s . Given our distributional assumptions, the average of the individual signals, that is, $\bar{y} = \sum_i y_i/N$ is normally distributed with mean s and precision defined as $\pi_s = (\sum_i b_i^2)(1 - \gamma)^2/\sigma_e^2$. Since s is a noisy signal of value, the precision of the specialist's information is $\pi_s = \pi\psi/\pi + \psi$. This is clearly increasing in the precision term π .

1.3.3 The information content of the limit order book. The specialist's ability to observe the individual orders constitutes a significant informational advantage. In particular, consider a trader who just observes the *aggregate* excess demand function $\sum q_i(p) = (\sum a_i) - (\sum b_i)p$, but not the individual orders themselves. This agent can form a statistic, denoted by y , where

$$y = \frac{\sum a_i (\sum b_i)^{-1} - \mu\gamma}{1 - \gamma} = s - \frac{\sum e_i}{\sum b_i(1 - \gamma)}. \quad (8)$$

This signal is normally distributed with mean s and precision

$$\pi' = \frac{(\sum_i b_i)^2 (1 - \gamma)^2}{N\sigma_e^2}. \quad (9)$$

Comparing π and π' , we find that the precision of the transparent system is always greater than the precision of the nontransparent system, provided that at least two traders differ in their risk tolerances.

The idea behind this result is straightforward. For example, an order to buy a large block of stock may be viewed very differently from, say, an order to sell an odd lot. In our model, an individual whose risk tolerance is high (manifested in the form of high price elasticity of demand) will not have a large portfolio hedging component to his or her trade. As a result, the specialist's statistical inference for this individual will be especially informative relative to others. Formally, if b_i is high, the precision of the statistic y_i is large. Simply observing aggregate demand does not allow the specialist to identify and filter out the more informative observations.

Transparency has another source of value as well. If $K > 0$, that is, if there are uninformed traders, the order form may convey information to the specialist regarding the asset's value. Formally, in a transparent system where the individual orders are observed, the specialist can form the statistic $\bar{y}$ as described above. In a system where only the aggregate demand function is observed, the specialist observes only the statistic y . Although y is an unbiased signal of value, its precision is larger because the imbalances that are uninformative cannot be distinguished from price-sensitive orders. Formally, the precision of the signal y is

$$\pi' = \frac{(\sum_i b_i)^2 (1 - \gamma)^2}{(N\sigma_e^2 + K\sigma_x^2)}, \quad (10)$$

which is strictly lower than before. It is for this reason, perhaps, that many call markets provide traders with information regarding imbalances arising from market orders. The arguments above demonstrate formally that there is a potentially large informational gain from observing the evolution of the limit order book prior to the opening.

1.3.4 The behavior of a strategic specialist. Having described the process by which the specialist's beliefs are determined, we turn now to the problem facing a strategic specialist who sets the opening price. Market clearing implies that at the price selected by the specialist, denoted by p_0 , aggregate excess demand is zero. From the definition of aggregate demand and the market-clearing condition $Q(p_0) + z = 0$, we obtain the opening price as a function of the specialist's trade,

$$p_0 = p^* + \lambda z, \quad (11)$$

where p^* is the market-clearing price in an automated call market and $\lambda = 1/(\sum_i b_i)$. Note that because traders' demands are conditioned on price (which conveys no useful information to them), the behavior of the specialist

is irrelevant and hence demands are unaffected by his or her presence.⁷ Given our assumptions, maximization of expected utility is equivalent to maximizing the certainty equivalent

$$v_s(z + e_s) + c_s - p_0(z)z - \delta(p_0(z) - p^c)^2 - \left(\frac{1}{2}\right)\rho_s\theta^2(z + e_s)^2, \quad (12)$$

where we explicitly recognize the dependence of the opening price on z . Differentiating this expression with respect to z and using Equation (11), the specialist's optimal trade can be written as

$$z = \Gamma_1(v_s - p^*) - \Gamma_2e_s - \Gamma_3(p^* - p^c), \quad (13)$$

where the coefficients Γ_1 , Γ_2 , and Γ_3 are positive constants given by $\Gamma_1 = 1/(2\lambda + \rho_s\theta^2 + 2\lambda^2\delta)$, $\Gamma_2 = \rho_s\theta^2\Gamma_1$, and $\Gamma_3 = 2\delta\lambda\Gamma_1$.

Using Equation (13), the close-to-open return, $p_0 - p^c$, can be rewritten as

$$p_0 - p^c = \lambda\Gamma_1(v_s - p^*) - \lambda\Gamma_2e_s + (1 - \lambda\Gamma_3)(p^* - p^c). \quad (14)$$

Equation (14) has a natural economic interpretation. The close-to-open return, which is based on the specialist's choice of opening price, has three components. The first component represents information known to the specialist about future prices, and is proportional to the difference between the specialist's estimate of value and the market-clearing price in the specialist's absence. The second term represents the effects of inventory control. Other things being equal, the specialist will move prices down to hedge a long opening inventory position and move prices up to hedge a short position. The third term captures the effects of exchange inducements to maintain price continuity. It is straightforward to show that the coefficient $1 - \lambda\Gamma_3$ lies in the interval (0,1) so that, other things being equal, the specialist will set an opening price that results in a close-to-open return that is a fraction of the close-to-open return in a fully automated system.

Equation (14) is a natural candidate for empirical estimation using a regression model. This would allow us to make direct inferences regarding the three motivations for specialist trading at the open identified by our theory. Estimation is, however, a nontrivial econometric problem, as we discuss later in Section 3 when we test the model.

⁷ In other words, the excess demand function's intercept on the price axis is unaffected by specialist's strategy. This is intuitive because the intercept is the mean reservation price which is unaffected by the specialist's presence.

1.3.5 The informational efficiency of prices. The model also sheds light on the informational efficiency of prices set by the specialist to the price relative to the price in an automated call market. To do this, it is useful to examine the pricing error, that is, the noise added by the trading process. We define the pricing error to be the deviation between the price and the value estimate based on the finest available information set, that is, that of informed traders. Observe that the (unconditional) variance of the pricing error in the automated call market without dealer intervention is

$$\text{var}[p^* - v_0] = \text{var}[\omega], \quad (15)$$

where ω is the noise term defined above. Similarly, the (unconditional) variance of the pricing error at the open is

$$\text{var}[p_0 - v_0] = \text{var}[\omega] + \lambda^2 \text{var}[z] + 2\lambda \text{cov}[z, \omega]. \quad (16)$$

This equation shows that the variance of the opening price relative to asset value is the sum of three terms: the noise around the Walrasian price, the variance of the specialist trade, and the covariance between the specialist's trade and the noise induced by noninformational trading. Comparing the variance of the two pricing errors, we obtain a measure of the informational inefficiency of the opening auction relative to the automated market:

$$\nabla \equiv \text{var}[p_0 - v_0] - \text{var}[p^* - v_0] = \lambda^2 \text{var}[z] + 2\lambda \text{cov}[z, \omega]. \quad (17)$$

The specialist's trade has positive unconditional variance arising from inventory hedging and price stabilization. This term adds to the inefficiency of the opening price relative to the Walrasian price. The second term, however, may mitigate this if the specialist is the contraparty to imbalances arising from noise traders, inducing a negative covariance. From a theoretical viewpoint, the difference in Equation (17) may be positive or negative. To see this, first consider the case where the specialist is risk neutral and the cost of maintaining price continuity is small so that the speculative motive is the only reason for trade. In this case, it is easy to show that the specialist's trade offsets, in part, the imbalance arising from liquidity motivated traders. Consequently, the specialist's actions will move price toward value, increasing informational efficiency. On the other hand, if inventory and continuity considerations are important but the specialist's information is poor, specialist trading is noninformation motivated. In this case, the addition of noise to the price discovery process produces less efficient prices. Thus the extent to which specialist trading facilitates the process of price discovery is an empirical one.

The model can also be extended to incorporate another realistic feature of the NYSE, namely competition from strategic floor traders or speculators of the type discussed by Werner (1998). Strategic floor traders will place buy (sell) orders if the price is expected to rise (fall) at the open relative to the

previous day's close. Such actions increase the profits of floor traders (at the expense of the specialist) because the specialist is selling (buying) to maintain price continuity, depressing (elevating) prices. So although specialists' enjoy considerable privileges at the open through their unique position and ability to set price, these factors do not necessarily give rise to economic rents.

1.4 Summary and discussion

The main findings from the model can be summarized as follows:

- There is a potentially large informational gain from observing the evolution of the limit order book prior to the opening. This gain comes from observing the individual demand functions of traders and also the aggregate imbalance arising from market-on-open orders.
- The opening price set by a strategic specialist reflects information obtained from observing the evolution of the limit book as well as noninformation factors such as inventory and price continuity. Specialist trades are positively related to the information component and negatively related to inventory and to the change from the previous day's close.
- If specialist inventory and stabilization trading is relatively small, the specialist's presence increases price efficiency.
- The specialist's informational and price-setting advantages need not result in positive expected profits because price stabilization ("leaning against the wind") is costly.

Before moving to the empirical analysis, it is helpful to discuss the model's underlying assumptions as far as they relate to the empirical analysis. A key assumption in the model is that traders can submit price schedules, obviating the need to forecast the specialist's pricing choice, as noted above. In reality, traders submit rectangular demand functions and thus may trade using market orders if they are sure that they want to be on a certain side of the market irrespective of the opening price. This is possible since the NYSE will allow for order cancellation or revisions if the price significantly departs from the previous day's close. If so, a proxy for p^* based on order-level data is most meaningful for opening prices within a range around the previous day's close. We discuss this point further in the empirical section.

Another key assumption is that traders act as price takers. In a small stock, where the number of traders placing orders is small, each trader has an influence on the specialist's beliefs and hence on price. In this case, traders may break up their orders or use combinations of limit and market orders to confound the specialist's inference problem. Thus our model should be interpreted as applying to active stocks where large numbers of traders participate at the open.

Finally, in our model the specialist's informational advantage came from the ability to observe incoming orders. However, that there are other avenues

by which the specialist might obtain information. Benveniste, Marcus, and Wilhelm (1992) argue that the specialist can make inferences about the likelihood of a trader having private information based on their identity or that of their broker. Further, the specialist can observe demand conditions in the market as a whole, as discussed in Amihud and Mendelson (1991) and Wohl and Kandel (1997). Specialists often cite factors peculiar to the floor, such as the ambient noise level, as important elements in their trading decisions. We can extend the model to incorporate information signals to the specialist arising from other than the order flow without altering our general conclusions.

2. Institutions and Data

2.1 Opening procedures on the NYSE

The theoretical model developed in the previous section corresponds closely to the actual institutional structure of the NYSE, but some subtleties of the exchange's protocols require more discussion. Trading on the NYSE takes place between 9:30 A.M. and 4:00 P.M. (Eastern Standard Time) Monday to Friday. After the opening, the NYSE operates as a continuous auction with a designated dealer or specialist who receives all market orders and maintains the public limit order book. As the center of trading on the exchange, the specialist also supervises the trading process, matches buyers and sellers, acts as an agent for other brokers, and exercises crowd control to ensure price and time priority and efficient order representation. Yet despite the specialist's prominent position on the exchange, most trading on the exchange is between public investors without specialist intermediation. Indeed, in recent years, the specialist participation rate (defined as total specialist share purchases and sales in all stocks divided by total NYSE share volume) has averaged about 18%.⁸

Opening protocols on the NYSE are different from the continuous trading system in many important respects. The NYSE's Opening Automated Report Service (OARS) stores the overnight accumulation of orders submitted electronically through the SuperDOT system. As orders are received, OARS continually matches ("pairs") buy and sell orders. In addition to system orders, floor brokers (the "crowd") who want to participate at the open give their preopening orders to the specialist who enters these into OARS. Floor brokers can give the specialist two types of orders for execution without being physically present: standard limit orders, and percentage ("participation") orders. Percentage orders allow the floor trader to participate in a selected percentage of the specialist's own trade. Sofianos and Werner (1997) report that most orders left with the specialist are percentage orders, suggesting that floor traders view specialist's opening trades as profitable.

⁸ Hasbrouck, Sofianos, and Sosebee (1993) provide a detailed description of NYSE systems and trading protocols; Madhavan and Sofianos (1997) analyze specialist participation across stocks and over time.

Floor brokers may also participate actively at the open by standing at the specialist post, in which case they might not necessarily submit their orders through OARS. Sofianos and Werner (1997) report, however, that the value of broker executed *active* trades at the opening (i.e., excluding OARS) for all stocks is just 0.9%, although the overall figure in continuous trading is 35%.

At the open, the specialist sets a *single* opening price at which the accumulated order imbalance from market-on-open and public limit orders must be absorbed by the crowd and the specialist's inventory. It is important to note that the specialist is not required to act as an auctioneer, but can actively trade for his own account. The specialist thus has considerable market power. This power is enhanced by the specialist's privileged access to the limit order book.⁹ In an illiquid stock, where the limit order book is especially thin, the specialist may open trading by posting a bid and offer price based on the limit order book or his own willingness to trade.

The specialist has an "affirmative obligation" to provide price continuity and maintain liquidity. Although there are no explicit penalties for violating this obligation, a specialist's repeated failure to adhere to these guidelines could ultimately be penalized in several forms, including reassignment of stocks, the failure to be allocated new, profitable stocks, or censure. The specialist can delay (subject to approval by a designated floor official) the opening or temporarily halt trading. Such a delay requires unusual circumstances, such as a "news pending" announcement or large imbalances. During this delay, nonbinding quote indications are usually issued on the tape to signal the source of the delay and to attract contraparty interest from the crowd.

2.2 Data sources and procedures

2.2.1 TORQ data. The data for this study are drawn from the Trades, Orders, Reports, and Quotes (TORQ) database that is publicly disseminated by the NYSE. The TORQ data include all trades, quotes, and system orders for a randomly selected set of 144 securities for the period November 1990–January 1991.¹⁰ System orders include all *orders* placed through the NYSE's automated trading system, SuperDOT. In addition, the database provides details on the *identity* of members behind each trade from the NYSE's audit trail data, the Equity Consolidated Audit Trade (CAUD) file. Specialist trades, however, are not directly identified, although they may be inferred, as discussed later. Of special interest, the TORQ data provide us with a complete representation of all orders (including public limit orders and market-on-close orders) entering the system prior to the opening, with fields indicating type and whether they represent buy or sell orders. Order information is vital for

⁹ Floor brokers may under NYSE rule 2115 (with the specialist's consent) view the limit order book, but it is not widely displayed, thereby conferring a substantial informational advantage to the specialist.

¹⁰ The four files that form the TORQ data are the Consolidated Transactions file (CT), the Consolidated Quotes file (CQ), the System Order Database file (SOD), and the Consolidated Audit Trail file (CD). For a detailed description of the TORQ database, see Hasbrouck (1992).

our study because the opening price will typically reflect the influence of all buy and sell orders, not just those that actually receive execution.

We construct the limit order book at the opening for each day for each stock, as in Kavejecz (1999). The procedure involves four stages: In stage one, we identify all limit orders that were submitted prior to the sample period. These orders either appeared as preopening orders on the first day of the sample period, or were “backfilled” using cancellations or trades that had no corresponding submissions in the sample period. In stage two, we identify all order submissions during each day of the sample period. In stages three and four, we record trades and cancellations for each day, respectively. Thus the limit order book at the opening includes orders that were submitted but not canceled or filled prior to the opening. Except for limit orders placed far away from current market prices, this procedure captures all other limit orders that enter the SuperDOT system.

2.2.2 System-clearing price. The NYSE floor official’s manual specifies that all preopening orders must be entered into the system prior to the determination of the opening price by the specialist. Consequently we can construct an analog to the Walrasian price, the *system-clearing price*. While this concept is clear, its implementation requires some care. We use the following algorithm to compute the system-clearing price:

- (i) For each stock and each day in the sample, all eligible preopen, market-on-open, and limit buy and sell orders are identified. Error or canceled orders are deleted, as are orders without volumes. Market buy orders are assigned an arbitrarily high price $\bar{p}$, while market sell orders are assigned a limit price of zero. We exclude days when the opening is delayed and/or trading halted by a floor official.
- (ii) For a price, p , in the discrete pricing grid $\mathbf{P} = \{0, d, 2d, \dots, \bar{p}\}$, where d is the minimum tick, the *cumulative demand* is defined to be the total buy-side volume (as indicated in the TORQ data) at price p or lower; for sells, *cumulative supply* is the total volume offered at p or higher. For every possible price $p \in \mathbf{P}$ we compute *excess demand*, $Q(p)$, defined as the difference between cumulative demand and supply at that price.
- (iii) A *clearing price* is a price at which absolute excess demand over all $p \in \mathbf{P}$ is minimized. Formally, p^* is a clearing price if $|Q(p^*)|$ is smallest for $p^* \in \{0, \dots, \bar{p}\}$. Let $\mathbf{C}$ denote the set of clearing prices.
- (iv) If $\mathbf{C}$ contains a single price p^* it is defined to be the *system-clearing price*. Otherwise, the system-clearing price is defined as follows: If $Q(p) < 0$ for all $p \in \mathbf{C}$, the system-clearing price is defined to be the lowest price contained in $\mathbf{C}$. If $Q(p) > 0$ for all $p \in \mathbf{C}$, the system-clearing price is the highest price in $\mathbf{C}$. If $Q(p_1) > 0$ and $Q(p_2) < 0$ for two prices in $\mathbf{C}$ (where p_1 is the highest price in

- C for which excess demand is positive and p_2 is the lowest price for which excess demand is negative), the system-clearing price is $(p_1 + p_2)/2$.
- (v) The system-clearing price is said to be *undefined* when there are no eligible preopen orders on either the buy side of the book, the sell side, or both, and/or the price defined in step (iv) above is either 0 or $\bar{p}$.

The procedure above yields a unique system-clearing price provided there are eligible preopening orders on both sides of the book to make this economically meaningful. Observe that the demand and supply schedules are step functions, so it is possible that there are multiple prices that qualify as clearing prices. Rule (iv) breaks ties by equating the system-clearing price with the lowest possible clearing price when there is excess supply and the highest clearing price when there is excess demand. In the unlikely event that there are two prices at which excess demand is exactly equal to excess supply, the clearing price is the average price. Similarly, rule (v) ensures that the system-clearing price, if it exists, is economically meaningful.




Figure 1

This figure shows the opening demand and supply schedules for Cray Research Inc. on November 28, 1990. Also shown in the figure is the market-clearing price using all orders, the previous day's closing price, and the midquote price at 3 P.M. the same day.

To better understand the nature of this procedure, it is useful to examine a specific example. Figure 1 shows the opening demand and supply schedules for Cray Research, Inc. on November 28, 1990, constructed using limit orders entering the SuperDOT system. Unlike the standard Marshallian cross, the supply and demand schedules are plotted as a function of price, so that the vertical difference between the two lines represents excess demand or supply at that particular price. The (unique) system-clearing price is \$26.875, since absolute excess demand is the smallest at this price. The system-clearing price is not necessarily the opening price, which is set by the specialist. In the example above, the specialist opened Cray Research at \$28.75 and took the excess supply into inventory.

Why did the specialist set the price so far above the market-clearing price? Our model suggests three factors. One possibility is that the specialist, concerned about price continuity, set the price close to the previous day's close of \$28.50, as shown in the figure. Second, the specialist may have had a short position relative to desired or target inventory and hence purchased shares for inventory reasons. Alternatively, the specialist may have set the price based on the information content of the limit book.

To see this, consider what we term the *limit-clearing price*. The limit-clearing price is defined to be the system-clearing price in the absence of market-on-open orders, that is, the price that would clear an electronic or automated market without specialist or floor trading based only on price-contingent limit orders. This construct allows us to make inferences about the relative information content of imbalances arising from market-on-open orders. In particular, if price-insensitive market-on-open orders are placed by noninformation motivated traders, the limit-clearing price will have greater informational content than the system-clearing price. Given the supply and demand schedules, the limit-clearing price in the example given above is \$28.75, exactly equal to the price set by the specialist. It is possible then that the specialist set the price above the system-clearing price because he or she viewed the order imbalances arising from market-on-open orders as uninformative.

In this particular case, however, the specialist's actions contributed significantly to price discovery. Specifically, the midquote at 3:00 P.M. on January 24 was \$28.75, equal to the opening price, and substantially different from the system-clearing price. We will examine whether this is an isolated instance or is typical of a more general pattern.

2.2.3 Discussion. Implicit in the computation of the system-clearing price is the hypothesis that traders would not alter their behavior if the specialist did not trade. This is reasonable for a system where traders submit limit orders prior to the open (and hence condition on price), but may not be the case in a transparent auction (e.g., the Paris Bourse) where the limit order book is displayed to traders. In the latter case, traders may strategically cancel

orders or delay placing orders as they observe the process of price formation, possibly yielding a different market-clearing price. This is possible on the NYSE only when there are large differences from the previous day's close, an issue we explicitly address in our empirical analysis.

While the NYSE is not a transparent market, the additional presence of floor traders who understand the specialist's trading strategy complicates matters. In particular, the system-clearing price is computed using all orders in OARS, and thus omits the influence of *active* floor broker orders.¹¹ From a theoretical perspective, speculative trades originating from active floor traders will change the market-clearing price from p^* to $p^* + \lambda[\sum_h q_h]$. Since active floor trades are not observed, the system-clearing price we compute is more representative of the price that would prevail in a pure or automated auction. Thus the system-clearing price is meaningful when interpreted as the price that would prevail in a nontransparent single-price auction, without trading from the exchange floor.

As a practical matter, Sofianos and Werner (1997) find that the magnitude of active floor broker trading is just 0.9% of value. However, for the smaller stocks, floor broker participation is higher, especially so in the bottom two size deciles where the corresponding figures are 16.9% and 26.8%, respectively. Thus our conclusions regarding specialist trading are most appropriate for the more active stocks, and henceforth our discussion is primarily focused on the more active deciles.

2.3 Identification of specialist trades

The system-clearing price defined above does not require any knowledge of specialist trades. However, some of our analyses require data on specialist transactions. We identify specialist trades in the TORQ data using the algorithm developed by Edwards (1999) and later refined by Panchapagesan (1999). The idea behind the algorithm is straightforward. The TORQ data include detailed information on the identity of traders in its audit file. Unlike the original audit trail data with the NYSE (CAUD), this information is only partially complete, as certain traders' identities, including those of the specialists, are left blank. This makes it possible to use the omission of a trader identity code to flag such transactions as possibly involving the specialist. Using filters based on prior knowledge of the CAUD file, and the NYSE's policies and procedures, we develop an algorithm to identify specialist trades. It should be emphasized that the algorithm for identifying specialist trades is quite distinct from the procedure used above to determine the system-clearing price.

¹¹ Passive orders, including participation orders, are already entered into OARS by the specialist and are hence incorporated into our analysis.

Table 1
Representation of a sample trade in the CAUD file

	Compared trade size		Badge		Account*		Source	
	Buy	Sell	Buyer	Seller	Buyer	Seller	Buyer	Seller
1	2000	2600	KP	0717	I	P	SuperDOT limit	Crowd
2	500		KP		I		SuperDOT Market	
3	100		SPEC		S		Crowd	

2.3.1 Audit-trail data. Before proceeding further, it is necessary to understand how audit information is presented in the TORQ database and in the CAUD. Hasbrouck, Sofianos, and Sosebee (1993, Table 1) present an example of how a trade is represented in the CAUD. A single “regular way” trade (5-day settlement) was reported to the Consolidated Tape System at 9:43 A.M. for 2600 shares at a price of $70\frac{7}{8}$. The audit trail information for this trade (only the relevant fields are shown) as it is represented in the CAUD and the TORQ database (with the original variable names in quotes) are shown in Tables 1 and 2.

The badge field identifies the broker representing the order. The account field indicates the transaction type and the description of the trader behind the order. The source field identifies the origin of the order. Compared trade size refers to the printed trade size (in CTS) that matches with clearance data provided by the National Securities Clearing Corporation. The key differences between the CAUD and the TORQ audit data are as follows: (1) the badge field is not provided in the TORQ data; (2) the specialist account type S is not included, while all other account types are included in the TORQ data; (3) crowd trades, including specialist trades, have no identifiable source (BTYPE or STYPE) in the TORQ data unless they trade with SuperDOT orders. In such cases, their source is L2 (for trades with limit orders) or D2 (for trades with market orders).

2.3.2 An algorithm for identifying specialist trades. The two data fields identifying the specialist in the original CAUD are the badge field and the

Table 2
Representation of a sample trade in the TORQ file

	Quantity		Account*		Source	
	Buy (BUYCOMP)	Sell (SELCOMP)	Buyer (BUYACCT)	Seller (SELACCT)	Buyer (BTYPE)	Seller (STYPE)
1	2000	2000	I	P	L1	L2
2	500	500	I	P	D1	D2
3	100	100		P		

*I represents nonprogram trading, nonmember, individual investor;
P represents nonprogram trading, member proprietary;
S represents specialist proprietary.

account type. As seen in the above example, only one of the two—the trader's account type—is available in the TORQ data. Using the account type and the order's source, we can identify potential specialist trades in the TORQ data.

To explain how the algorithm works, consider trades where the specialist could have been a possible buyer. Since the problem of identification is symmetric, the description applies to specialist sales as well. The algorithm identifies only the proprietary trades of the NYSE specialist at the NYSE.¹² NYSE rule 132 mandates the provision of account information for audit trail purposes by all traders. Therefore, account types cannot be missing in the audit data unless they were systematically excluded. Since the specialist's account type (S) is missing in the TORQ data, the necessary condition for a specialist buy is that the buyer account type (BUYACCT) should be blank. Additional refinements are needed because account types can be missing for nonspecialist trades as well. These are discussed in detail in Panchapagesan (1999). Formally, the algorithm identifies specialist buys (sells are symmetric) in the TORQ audit data as records where (1) account type (BUYACCT) is missing, (2) source (BTYPE) is D2, L2, or blank, and (3) any intermarket trading system (ITS) trades (BTYPE = I1 or I2) must not have records in the system order database (SOD) file included with TORQ.

To the extent that the TORQ database was constructed in the manner conjectured, the algorithm should be highly accurate since the data are extracted from the NYSE's own audit-trail records.¹³ Panchapagesan (1999) examines three different sources that use validated specialist trade data to benchmark the algorithm. In particular, the algorithm yields a volume-weighted specialist participation rate of 16.3% for the 144 TORQ stocks as compared to 19.8% for all NYSE stocks, as reported in the NYSE *Fact Book* (1991). Panchapagesan (1999) shows that the implied participation rate in TORQ stocks is similar to that reported by Madhavan and Sofianos (1997) and demonstrates the same patterns across stocks. Additional specialist trading statistics such as the stabilization ratio and price position of specialist trades are also compared and shown to be similar to those reported in the NYSE *Fact Book* (1991) and Sofianos (1995). We use this procedure to identify all specialist transactions, including the open, for all stocks in the TORQ data. We compute a constructed opening share inventory for each day and each stock as the sum of all signed specialist trades. It should be noted that the specialist's actual inventory level at the start of the sample is not observed, so that the constructed inventory is correct up to an unknown constant.

¹² The audit information is incomplete for trades reported by other exchanges. This algorithm will therefore not capture trades of NYSE specialists routed through other exchanges. However, such trades are likely to constitute a low percentage of their total trades.

¹³ Conversations with exchange officials suggest that this is the case.

3. Empirical Results

3.1 Is an automated call market feasible at the open?

3.1.1 Descriptive statistics. Table 3 presents descriptive statistics on all NYSE-listed stocks in the TORQ database and for deciles of dollar trading volume. It is clear that the sample stocks vary widely in important dimensions including firm size, trading frequency, and price. Size and trading frequency increase monotonically with deciles of dollar trading volume. Overall, the sample stocks are representative of the entire universe of NYSE stocks in these dimensions. We also report the percentage of daily dollar trading volume at the call auction open and during the first half-hour of trading from 9:30–10:00 A.M. Over all stock-days, 17.5% of the value of trading takes place in the opening half-hour. The fraction of daily volume at the open and in the first-half hour is inversely related to trading activity, indicating that the open is an especially important period for small stocks.

3.1.2 Operation of the NYSE opening. Table 4 presents a more in-depth analysis of the operation of the NYSE open for all stocks and for deciles of dollar trading volume. A call or batch market mechanism is used to open the market in 67% over all stock-days in the sample. In the remaining 33%, the market opens with a two-sided quotation representing public limit orders or the specialist's own willingness to trade. However, as shown in the table, the reliance on batch mechanisms monotonically *increases* with trading activity. In the top decile, a call market is used in virtually all cases, while in the

Table 3
Summary statistics on NYSE openings

	Market capitalization (in \$ million)	Average number of trades a day	Average price (\$)	Average daily dollar trading volume		
				During the day (\$ '000s)	At the open (%)	First half-hour (%)
All stocks	2721	52	19.9	4889	9.7	17.5
Deciles						
10 (High)	21,859	297	48.9	40,158	5.4	15.3
9	2311	57	28.3	3872	5.4	13.2
8	801	53	18.7	1709	5.5	13.8
7	451	32	22.3	983	7.0	14.5
6	329	22	26.0	469	8.6	15.1
5	346	16	16.4	241	11.1	17.9
4	152	11	14.9	115	15.2	19.8
3	61	9	10.4	59	20.6	24.2
2	29	6	4.7	24	23.0	26.7
1 (Low)	6	4	1.2	4	25.8	36.3

This table presents summary statistics for stocks in the TORQ database for the period November 1990–January 1991. Market capitalization is computed as of October 31, 1990, using CRSP data. All trading statistics are computed using NYSE trades only. Trading volume at the open refers to volume at the call market held by the specialist. The first half-hour volume includes all trades between 9:30 A.M. and 10:00 A.M. Summary statistics are computed for all sample stocks and for dollar trading volume deciles. Each decile average represents the simple mean of the daily averages of stocks within the decile.

Table 4
An analysis of NYSE opening protocols

	Days with a batch opening (%)	Days with de- layed opening (%)	Days when the system-clearing price was defined (%)	Days when spe- cialist traded at the open (%)	Mean share im- balance as a ratio of opening share volume (%)
All stocks	66.7	0.8	92.1	73.4	40.8
Deciles					
10 (High)	99.3	1.3	95.3	60.1	34.5
9	94.2	0.5	88.2	61.9	42.9
8	87.3	0.9	93.3	64.7	46.0
7	75.6	1.1	87.3	73.2	44.7
6	69.4	0.6	88.8	68.5	40.1
5	59.3	0.8	93.8	72.4	47.8
4	55.5	0.4	95.0	83.8	38.2
3	46.1	0.4	89.9	84.0	33.7
2	41.8	1.6	98.6	83.1	33.8
1 (Low)	20.4	0.6	91.1	89.4	39.4

This table presents summary statistics on system-clearing prices at the NYSE open for all stocks in the TORQ database (excluding ticker symbols UTD and MBK) and for deciles of dollar trading volume. In computing the absolute excess demand at the system-clearing price (last column), we use only days with a call market at the open. The system-clearing price is the price that clears the market if floor participants (including the specialist) do not trade. It is undefined when there are no eligible preopen orders (1) on the buy side, (2) on the sell side, or (3) on both sides of the book. The specialist trade percentage (column 5) is only computed for days with an opening batch auction. Each decile cell represents the simple averages within that decile.

smallest decile, it is used on only 20% of stock-days. Table 4 also reports the frequency with which the specialist delayed the quote. Delays are relatively infrequent (0.8% of all stock-days) and do not appear to be systematically related to trading activity. Occurrences are equally split between so-called news pending announcements and halts caused by imbalances. Although these events are interesting, the frequency of delayed openings in our sample is too small to permit a meaningful statistical analysis of the open at these times. Since our focus is on the operation of the call market on “normal” trading days, in what follows we restrict attention to openings that are not delayed or halted.

3.1.3 Dealer versus batch mechanisms. Table 4 also reports the frequency with which the system-clearing price at the NYSE open is defined for all stocks and for dollar trading volume deciles in the TORQ database. Recall that the system-clearing price clears the market if floor participants (including the specialist) do not trade. It is undefined when there are no eligible preopen orders on the buy and/or sell sides of the book. Of special interest, the system-clearing price is defined in 92% of all stock-days, although a call market is actually used in only 67% of stock-days. There is no systematic pattern in the viability of a pure auction that is discernible across deciles of trading activity. These results suggest that an automated call market (without dealer intervention) could function as a viable alternative to the existing NYSE protocols for most active stocks, although with possibly large price volatility as discussed below. The percentage of stock-days (with call auctions) when the specialist

traded at the open is inversely related to trading activity in the stock. The average volume in market-on-open orders (defined as the market-on-open buy volume plus market-on-open sell volume, divided by two) is approximately 40% opening volume.

3.2 Dealer trading behavior

3.2.1 An econometric model of price formation at the open. Our theoretical model demonstrates that the specialist's optimal action at opening is a function of three factors: the information provided by the limit order book, the specialist's inventory position, and exchange-mandated price continuity rules. In this section, we develop an econometric model to test these predictions and hence assess the net effects of dealer actions at the open.

Although the inventory and price continuity factors can readily be estimated, econometric tests of dealer behavior are complicated because we do not observe the dealer's estimate of fundamental value. It is tempting to simply use an estimate of future value in Equation (13) as a proxy for the specialist's conditional expectation and estimate the model using ordinary least squares. However, this could induce an errors-in-variables problem that might bias all the OLS coefficient estimates. Fortunately, there is a solution to this problem in the form of instrumental variables estimation.

Let $\mathbf{X}$ denote a vector of state variables (e.g., include market order imbalances at the open, the previous day's close, and the system-clearing price) observed by the specialist *prior* to determining the opening price. Using a linear specification for the conditional mean together with Equation (14) and we obtain the system:

$$E[v - p^* | \Omega_s] = \gamma' \mathbf{X} \quad (18)$$

$$p_0 - p^c = \beta_0 + \beta_1(E[v - p^* | \Omega_s]) + \beta_2 e_s + \beta_3(p^* - p^c). \quad (19)$$

The equation system [Equations (18) and (19)] can be jointly estimated using an instrumental variables (IV) estimator. Theory indicates the coefficients in Equation (19) have the following signs: $\beta_1 > 0$, $\beta_2 < 0$, and $0 < \beta_3 < 1$.

3.2.2 Model estimates. Table 5 presents the limited-information maximum likelihood (k -class) IV technique estimates of the model for deciles of dollar trading volume for stocks in the TORQ database.¹⁴ Figures in parentheses represent asymptotic standard errors computed from the Fisher information matrix. The independent state variables in Equation (18) are the total market-on-open order imbalance in shares (difference between buy and sell volume relative to average daily share volume), the system-clearing price, and the previous day's close. We use the same-day 3:00 P.M. midquote price as

¹⁴ The k -class estimators are efficient under assumptions of normality.

Table 5

Transaction-level analysis of opening returns

	β_0	β_1	β_2	β_3	No. of observations
10 (highest decile)	-0.00 (0.00)	0.15 (0.02)*	0.04 (0.09)	0.30 (0.02)*	782
9	-0.00 (0.00)	0.16 (0.02)*	0.00 (0.04)	0.20 (0.01)*	716
8	0.00 (0.00)	0.09 (0.02)*	0.21 (0.09)*	0.16 (0.01)*	667
7	0.00 (0.00)	0.07 (0.02)*	-0.02 (0.08)	0.15 (0.01)*	545
6	0.00 (0.00)	0.10 (0.02)*	0.02 (0.04)	0.19 (0.02)*	497
5	0.00 (0.00)*	0.14 (0.02)*	0.03 (0.05)	0.24 (0.02)*	499
4	-0.00 (0.00)	0.14 (0.02)*	0.02 (0.03)	0.21 (0.02)*	430
3	-0.00 (0.00)*	0.16 (0.02)*	0.09 (0.11)	0.38 (0.03)*	280
2	-0.00 (0.00)*	0.20 (0.04)*	-0.17 (0.11)	0.32 (0.03)*	310
1 (lowest decile)	-0.03 (0.01)*	0.73 (0.05)*	-0.14 (0.26)	0.87 (0.04)*	96

This table presents limited-information maximum likelihood estimates of a joint model of specialist beliefs and returns at the open for dollar trading volume deciles in the TORQ database (excluding ticker symbols UTD and MBK). The model is

$$p_0 - p^c = \beta_0 + \beta_1(E[v - p^*|\Omega_s]) + \beta_2\epsilon_s + \beta_3(p^* - p^c),$$

where $E[v - p^*|\Omega_s] = \gamma'X$. Future value, v , is the (log) midquote at 3:00 P.M. The instruments in X are the market order imbalance at the open (buy volume less sell volume) as a percentage of average daily volume for the stock, the (log) system-clearing price, and the (log) previous day's closing price. The variables in the opening return equation are the (log) opening price set by the specialist, p_0 , the (log) previous day's closing price, p^c , the (log) system-clearing price, p^* , and the specialist preopening inventory as a percentage of average daily volume for the stock, ϵ_s . Figures in parentheses are asymptotic standard errors. An asterisk denotes significance at 5% level.

the benchmark future value. We selected the 3:00 P.M. midquote because of the possibility that the specialist may “validate” his or her choice of the opening price over short horizons. Garbade and Sekaran (1981) discuss this issue, but conclude it “appears unlikely that a specialist can ‘peg’ the market in an active issue over any appreciable length of time, such as 30 minutes.” The low specialist participation rates for active NYSE stocks [see, e.g., Madhavan and Sofianos (1997)] also suggests that validation is unlikely to be a significant factor for our benchmarks. Indeed, our results also hold for other benchmarks, including the 10:00 A.M. midquote and the same day closing price. In Equation (19), the specialist's opening inventory is expressed as a percentage of inventory over the average daily share volume for the stock. Finally, we scale the price variables by taking log transformations. The average R^2 across deciles is 66% for the conditional mean regression and 33% for the close-to-open return regression, suggesting that the model works well for the given data.

The coefficient on the specialist's expected return is positive and significant for all deciles. This suggests that the specialist sets prices that move the opening price in the direction of his value estimate. These results are robust to the choice of instruments. However, the coefficient on specialist inventory is

not significant, except for one decile where it is of the wrong sign. The weak inventory results parallel previous studies of specialist behavior using inventory data [see, e.g., Hasbrouck and Sofianos (1993), Madhavan and Smidt (1993), and Madhavan and Sofianos (1997)] that fail to detect evidence for inventory control during the day.¹⁵ Finally, there is evidence that specialists engage in price stabilization. The coefficient on the return from the previous day's close to the open without specialist participation lies between zero and one as predicted and statistically significant in all deciles. The specialist will open at prices above (below) the previous day's close if the stock will open above (below) the previous day's close without his or her participation. But the extent to which the opening price is set away from the previous close is of a smaller magnitude than what would have resulted if he or she were not present, as the coefficient is less than one.

In other words, the specialist ensures that opening prices reflect market-clearing prices, although not to its full extent. It may provide profit opportunities for agents using very short-horizon technical or momentum strategies. In an efficient market, such stabilization is likely to be costly, and the fact that specialists are willing to do so suggests that the reputational costs of failing to comply with the exchange's affirmative obligations are perceived as large.

3.3 Do dealers speed price discovery?

3.3.1 The relative contributions to price discovery. The IV estimates above provide strong evidence that specialist trades reflect an information factor. But dealer trades also reflect noninformational factors including price continuity and inventory control that add noise to the opening so the net effect of specialist trading on price discovery is not immediately obvious.

We begin with a direct test of the relative predictive power of the opening and system-clearing prices that is not predicated on a particular model of specialist behavior. Specifically we are interested in whether the price set by the specialist is a better estimate of future value than the price if specialist trading were constrained to zero. Table 6 presents results of a two-stage regression analysis of opening prices on the midquote at 3 P.M. on the same day:

$$\ln(p_{i,t}^+ / p_{i,t-1}^c) = \alpha_i + \beta_i \ln(p_{i,t}^1 / p_{i,t-1}^c) + e_{i,t} \quad (20)$$

$$\hat{e}_{i,t} = \gamma_i + \delta_i \ln(p_{i,t}^2 / p_{i,t-1}^c) + u_{i,t} \quad (21)$$

¹⁵ As the specialist's inventory at the start of the sample period is not observed, the intercept captures this as well as any desired (or target) inventory positions the specialist maintains. However, the intercepts are not significantly different from zero, possibly because desired inventories are close to zero or because the time series is relatively short. It is also possible that there are adjustments to inventory (including odd-lot trades, error corrections, off-exchange trades, and transfers after hours) that are not captured by our constructed inventory variable.

Table 6
Informativeness of opening prices

	$\hat{\alpha}_i$	First-stage regression				Second-stage regression			
		$\hat{\beta}_i$	R^2	% stocks with $\hat{\beta}_i \neq 1$	$\hat{\gamma}_i$	$\hat{\delta}_i$	R^2	% stocks with $\hat{\delta}_i \neq 0$	
Panel A:		Specialist opening price				System-clearing price			
All stocks	0.00	0.88	0.27	22.7	−0.00	0.03	0.03	10.2	
Decile 10	0.00	0.89	0.27	14.3	0.00	0.04	0.02	7.1	
Panel B:		System-clearing price				Specialist opening price			
All stocks	0.00	0.18	0.12	96.6	−0.00	0.64	0.15	63.6	
Decile 10	0.00	0.24	0.14	100.0	0.00	0.59	0.13	85.7	
Panel C:		Specialist opening price				Limit-clearing price			
All stocks	0.00	0.87	0.27	20.8	0.00	0.06	0.03	11.5	
Decile 10	0.00	0.86	0.29	14.3	−0.00	0.01	0.00	0.0	
Panel D:		Limit-clearing price				Specialist opening price			
All stocks	0.00	0.26	0.10	88.5	−0.00	0.71	0.19	78.1	
Decile 10	0.00	0.53	0.11	64.3	0.00	0.56	0.14	78.6	

This table presents results of a two-stage regression analysis for the model:

$$\ln(p_{i,t}^+/p_{i,t-1}^c) = \alpha_i + \beta_i \ln(p_{i,t}^1/p_{i,t-1}^c) + e_{i,t}$$

$$\hat{e}_{i,t} = \gamma_i + \delta_i \ln(p_{i,t}^2/p_{i,t-1}^c) + u_{i,t},$$

where for stock i and day t , $p_{i,t}^+$ is the same day 3 P.M. midquote price, $p_{i,t-1}^c$ is the close on the previous day, $p_{i,t}^1$ and $p_{i,t}^2$ represent two possible prices (the system-clearing price or the opening price set by the specialist), and $\hat{e}_{i,t}$ is the estimated residual from the first-stage regression. The regressions are estimated individually for all stocks (excluding ticker symbols UTD and MBK) where the system-clearing price is defined for at least 30 days. Each cell represents the mean value of estimates across all individual regressions. We report the percentage of stocks for which $\hat{\beta}_i \neq 1$ and $\hat{\delta}_i \neq 0$ at the 5% significance level in a two-tailed test. Panels C and D replace the system-clearing price with the limit-clearing price.

where, for stock i and day t , $p_{i,t}^+$ is the future price, $p_{i,t-1}^c$ is the close on the previous day, $p_{i,t}^1$ and $p_{i,t}^2$, represent two possible prices (either the system-clearing price or the opening price set by the specialist), and $\hat{e}_{i,t}$ is the estimated residual from the first-stage regression.

The idea behind this approach is straightforward. If the price $p_{i,t}^1$ is a good predictor of future value in Equation (20), then price $p_{i,t}^2$ should have no explanatory power with respect to the estimated residuals in Equation (21). If this is the case, when we use $p_{i,t}^2$ as the independent variable in Equation (20) and $p_{i,t}^1$ in Equation (21), the estimated coefficient of $p_{i,t}^2$ should be nonzero. Accordingly, we estimate the model twice, reversing the order of the two prices used as independent variables. Later, we use the limit-clearing price instead of the system-clearing price to differentiate the information content of market-on-open orders from limit orders. We estimate the models in return form to avoid econometric problems induced by possible nonstationarity in price levels using the same-day 3:00 P.M. midquote as our benchmark of future value.

Table 6 summarizes the results of this estimation. We estimate the model for the 89 stocks for which the system-clearing price is defined on 30 or more days; the estimates without this filter are, however, very similar and are not reported. We report mean coefficient estimates and the mean R^2 from individual regressions using a portfolio of all 89 stocks and a portfolio of 14 stocks comprising the most active stock decile by dollar trading volume. We report the percentage of stocks for which $\hat{\beta}_i \neq 1$ in the first stage and $\hat{\delta}_i \neq 0$ in the second stage at the 5% significance level in a two-tailed test. In a naïve random walk model $\beta = 1$, although in our model $\beta < 1$ because transitory order imbalances at the open add noise to prices. We do not, however, test a particular model of returns but rather focus on δ to capture the incremental explanatory power of the second price, respectively.

Panels A and B use the system-clearing price, while panels C and D use the limit-clearing price as the market-clearing price. Consider first panel A. In the first stage, using the specialist opening return as the independent variable, the model fits well, and in the second stage, the system-clearing return has little explanatory power over the residuals from the first-stage estimation. Only 10.2% of stocks had $\hat{\delta}_i$ significantly different from zero in the second stage.

The converse, however, is not true. When the system return is used in the first stage as the independent variable fit is poor and more importantly, residuals from the second-stage regression are explained by the opening return.¹⁶ Indeed, the coefficient δ is significantly different from zero for 63.6% of stocks. Though we report the percentage of stocks with $\beta \neq 1$ under a two-tailed test, we find that β is generally less than 1 as predicted by the model. The results are quite similar when we use the limit-clearing price, and are robust to the use of other benchmark prices including the same-day 10:00 A.M. midquote and the closing price. Thus the opening price set by the specialist has significant predictive power beyond the system-clearing price, but that the reverse is not the case.

3.3.2 Variance ratio tests. A related question concerns the distribution of opening and system-clearing prices, that is, price efficiency. Recall that in the market-clearing price $p^* = v_0 + \omega$, where ω is a stochastic disturbance term (*system pricing error*) with mean zero. Similarly, we can also write the opening price as $p_0 = v_0 + \varepsilon$, where ε is a noise term (*opening pricing error*) with an unconditional mean of zero. The question is then whether the variance of ω exceeds that of ε ; we can answer this question by comparing the variance of the opening price around a suitable future price to the variance of the system-clearing price around the same benchmark price.

Denote by σ_W^2 and σ_0^2 the variance of the system pricing (Walrasian) and opening pricing errors, that is, $\text{var}[\ln(p^+) - \ln(p^*)]$ and $\text{var}[\ln(p^+) - \ln(p_0)]$,

¹⁶ The left-hand side of the panels A and C are different in Table 6 due to the slight difference in the number of observations (stock-days) used. The system-clearing price and limit-clearing price were not defined on all stock-days.

where p^+ denotes the appropriate future price, p_0 is the opening price, and p^* is the system-clearing price. For days with no opening call market, p_0 is the opening midquote. We report results using the midquote at 3:00 P.M. on the same day as our estimate of future value. The results are robust to the use of other prices including the same-day 10:00 A.M. midquote and the closing price.

To begin with, consider the conventional variance ratio σ_W^2/σ_0^2 , defined as

$$\Lambda = \frac{\text{var}[\ln(p^+) - \ln(p^*)]}{\text{var}[\ln(p^+) - \ln(p_0)]}. \quad (22)$$

The average statistic for all stock-days for which there were call opens is 8.16, which is economically large. We computed the percentage of stocks for which the variance of the system pricing errors σ_W^2 is significantly greater than or less than the corresponding variance of the opening pricing error σ_0^2 in a one-sided test. The percentage of days for which $\sigma_W^2 > \sigma_0^2$ is statistically significant at 80.4%. By contrast, the percentage of stock-days for which $\sigma_W^2 < \sigma_0^2$ was statistically significant was zero. Similar findings hold for the individual deciles. Thus the variance of the system-clearing price is both economically and statistically much greater than the variance of the opening price.

Conventional variance ratio tests could suffer from econometric problems arising from overlapping observations, cross-stock correlations, and serial correlation and heteroscedasticity in the return series. Ronen (1997) proposes a generalized method of moments (GMM) test that addresses these problems. Ronen's approach was developed to test the properties of return variances over trading and nontrading periods [see also George and Hwang (1997)], but the GMM methodology can be directly applied to this problem.

Specifically, given N stocks we form $2N \times 1$ vector:

$$g(\Lambda) = \frac{1}{T} \sum_{t=1}^T \begin{pmatrix} R_{1,w}^2 - \Lambda_1 v_{1,o} \\ \vdots \\ R_{N,w}^2 - \Lambda_N v_{N,o} \\ R_{1,o}^2 - v_{1,o} \\ \vdots \\ R_{N,o}^2 - v_{N,o} \end{pmatrix}, \quad (23)$$

where $R_{i,w}$ denotes the system-clearing pricing error, $R_{i,o}$ is the opening pricing error, $v_{i,o}$ is the true variance of the opening pricing error, and Λ_i is the variance ratio of the pricing errors ($v_{i,w}/v_{i,o}$) for stock $i = 1, \dots, N$. The advantage of using GMM to estimate a just-identified system such as above is that the procedure incorporates covariances across ratio estimates that are often ignored in conventional variance ratio tests. We use the Wald statistic, $T(\hat{\Lambda} - 1)' \Sigma_{\hat{\Lambda}}^{-1} (\hat{\Lambda} - 1) \approx \chi^2(N)$ to test the null hypothesis $H_0: \Lambda =$

Table 7
Wald tests of price efficiency

	Variance ratio (σ_W^2/σ_0^2)					
	Using system-clearing prices (Wald statistic)			Using limit-clearing prices (Wald Statistic)		
	Using p^*	Using p_1^*	Using p_2^*	Using p^*	Using p_1^*	Using p_2^*
All stocks	8.49 (152.7)*	4.89 (206.1)*	2.27 (159.5)*	6.52 (115.9)*	3.79 (265.9)*	2.21 (222.1)*
Decile 10	10.75 (36.9)*	2.42 (166.5)*	3.19 (47.3)*	2.60 (34.1)*	1.57 (45.6)*	1.93 (35.9)*

This table presents generalized method of moments (GMM) estimates of variance ratios of the pricing errors associated with the opening price and system-clearing prices in call market opens for all stocks in the TORQ database and for the top decile (10) of stocks by dollar trading volume. The table also reports the associated value of the Wald test that the variance ratios for all stocks are jointly equal to one. Opening pricing error refers to the difference between the (log) future value and the (log) opening price; the system pricing error refers to the difference between (log) future value and the (log) system (or limit) clearing price. Future value is the midquote at 3:00 P.M. on the same day of the opening trade. We report results for three different subsamples of system-clearing prices: (1) using all stock-days (p^*), (2) using stock-days when the system-clearing price is within \$3 from the previous closing price (p_1^*), and (3) using stock-days when the system-clearing price is within 10% of the previous closing price (p_2^*). We denote by σ_W^2 and σ_0^2 the variance of the system pricing (Walrasian) error using all orders, and opening pricing errors, respectively. An asterisk denoted significance at the 1% level.

$[1, \dots, 1]'$, that is, the vector of return variance ratios is jointly equal to unity. Note that the Wald statistic explicitly uses the variance-covariance matrix of the variance ratio estimates (Σ_Λ). Ronen (1997) illustrates how omission of covariance terms can often bias the test statistic against the null. The advantage of a Wald test over other tests (e.g., the Lagrange multiplier test) is that we compute individual variance ratios for *each* stock while *jointly* testing the null hypothesis.

Table 7 summarizes the results of the GMM estimation for all stocks and for the top decile by dollar trading volume. For all stocks, the estimated variance ratio is 8.49 using the system-clearing price with a Wald statistic of 152.7, rejecting the null hypothesis at the 1% level.¹⁷ Using the limit-clearing price, the corresponding ratio for all stocks is only 6.52 (with a Wald statistic of 115.9). While the regression results of Table 6 indicate that point estimates of future returns are more accurate using all orders, these results indicate that the system-clearing price is less noisy if market-on-open orders are excluded. Together, these results suggest that there is both noise and information in market-on-open orders. The figures for the top decile of stocks exhibit a similar pattern.¹⁸ We obtain similar results focusing only on quote openings where batch markets are defined, again confirming our results.¹⁹

¹⁷ We use only stock-days when there was a call market and the system-clearing price is meaningfully defined and treat each decile as an individual portfolio.

¹⁸ Less active stocks have higher variance ratios and conventional tests easily reject the null hypothesis.

¹⁹ The ratios for the system and limit-clearing prices in quote openings are 15.1 and 11.3 with Wald statistics of 80.0 and 115.4, respectively.

As discussed earlier, specialists provide indicated prices to floor participants and regional specialists when the market-clearing price differs substantially from the previous day's closing price. In this case, strategic order cancellation could make our definition of a market-clearing price essentially meaningless. To include such a possibility, we redo the variance ratios using only those stock-days when the system-clearing price is within a specified range from the previous day's closing price. We consider two values for the range: (within 10% and within \$3 of the previous day's close. Focusing, for example, on the 10% range, the variance ratio for all stocks is 2.27 using the system-clearing price and 2.21 with the limit-clearing price. These ratios are significantly smaller than the corresponding unconditional ratios, as expected, but in all cases we nonetheless reject the null hypothesis at the 1% level. These results suggest that our conclusions are robust to other, narrower definitions of the market-clearing price.

3.4 Profitability of specialist opening trades

To assess the economic value of the specialist's position, we compute ex post returns for buys and sells. Anecdotal evidence indicates that many specialists have a relatively short holding periods and often liquidate their opening positions in the first few hours of trading, so we compute returns using a variety of within-day benchmark prices including the 10:00 A.M. midquote, the 3:00 P.M. midquote, and the closing price. The results were similar across benchmarks, and we report the results using the 10 A.M. midquote.

Specialist opening purchases (sales) are associated with positive (negative) returns (measured as log price difference between the 10 A.M. midquote price and the opening price) for almost all stocks and across all deciles of trading volume. In particular, the mean ex post return after specialist purchases was 0.24% (with a standard error of 0.04%), while after specialist sales it was -0.19% (with a standard error of 0.04%). The return figures were similar for the other two benchmarks, but were not significantly different from zero at the 5% significance level when using dollar and share volume weighting schemes. These returns are consistent with the evidence reported by Sofianos (1995) who estimates specialist gross trading profits using NYSE audit-trail data.

The results show that despite the specialist's informational advantages—identified in the IV estimation—gross trading profits are relatively modest. There are several possible explanations for this finding. First, price stabilization is costly since it implies some element of “leaning against the wind.” Second, the specialist faces direct competition from floor traders and indirect competition in the form of participation orders, as described by Sofianos and Werner (1997). Third, risk aversion might limit the extent to which the specialist is willing to take large positions based on information. Finally, specialists concerned about exchange scrutiny may avoid trading aggressively on their informational advantages for fear of creating the appearance of impropriety.

4. Conclusions

This article analyzes the NYSE opening with the objective of better understanding the process of price formation. We develop a model of a single-price auction and show the presence of a dealer can speed price discovery. Intuitively the dealer extracts valuable information from observing the evolution of the limit order book, and trades on this signal, increasing price efficiency. However, noninformational dealer trading (e.g., inventory rebalancing or price stabilization) impedes price formation so the ultimate impact on efficiency is an empirical issue.

We test the model using order-level data and obtain several interesting results. First, the opening price reflects specialist information and price continuity requirements. Although the specialist's information-based trading enhances price efficiency, price stabilization induces staleness by tying the opening price to the previous day's close. This raises the natural question as to why exchange-mandated price continuity rules exist in the first place. Dutta and Madhavan (1995) argue theoretically that price stabilization rules arise because they limit the ability of dealers to extract rents from their superior information. Price continuity rules also provide some assurance that the investor's order will execute at a price near that of the previous day's close, a datum almost certain to be known at the time of order submission. In the event that the opening price will depart substantially from the previous day's close, the specialist is required to put out "indicated" prices and allow time for off-floor traders to revise their orders, mitigating their informational disadvantage. Indeed, although specialist opening trades have positive returns, they are modest and are not significantly different from zero when volume weighted.

Second, and consistent with previous empirical research on intraday trading, inventory effects at the open are very weak. This is surprising given the specialist's unique advantages at the open in terms of their privileged access to the limit order book and market power. This finding provides additional evidence that dealers do not use price as a tool to control excess inventory. Just how dealers manage their inventories is an interesting topic for future research.

Third, there is strong evidence that the NYSE's designated dealer (specialist) sets a more efficient price than the price that would prevail in a pure call market using only public orders. It is important to recognize that this result is obtained in the context of existing protocols and does not imply that an opening with a designated dealer is preferred to a fully transparent call market where *all* investors observe the limit order book in the preopening period. Empirical studies are ultimately limited in examining such "what if" questions, but experimental methods [see, e.g., Bloomfield and O'Hara (1999), and Schnitzlein (1996)] offer considerable promise in this regard.

Fourth, system-clearing prices are not always defined and are highly sensitive to market-on-open order imbalances, especially in thinly traded stocks.

This result is important because of the widespread belief that single-price auctions are efficient mechanisms to aggregate diverse information. Economides and Schwartz (1995) and Schwartz and Wood (1995) argue for the wide use of call auctions when return volatility is high, such as after overnight nontrading periods, trading halts, or days when derivatives such as options or futures expire. Our results indicate that it is not practical to trade all securities—especially thinly traded assets—in a pure call market because relatively small imbalances can generate prices that are not economically meaningful. Indeed, reliance on the call auction actually increases with trading activity.

The process by which investors' latent demands are translated into realized prices and volumes is a highly complex process that we are only now starting to understand. Our results add to a growing body of evidence that highlights the crucial roles of information and market structure in determining price efficiency, but there are still many important questions to be answered before we fully understand the inner workings of the black box of trading mechanisms.

References

- Amihud, Y., and H. Mendelson, 1987, "Trading Mechanisms and Stock Returns: An Empirical Investigation," *Journal of Finance*, 42, 533–553.
- Amihud, Y., and H. Mendelson, 1991, "Volatility, Efficiency and Trading: Evidence from the Japanese Stock Market," *Journal of Finance*, 46, 1765–1790.
- Benveniste, L. M., A. J. Marcus, and W. J. Wilhelm, 1992, "What's Special About the Specialist? Floor Exchange Versus Computerized Market Mechanisms," *Journal of Finance*, 32, 61–86.
- Biais, Bruno, P. Hillion, and C. Spatt, 1999, "Price Discovery and Learning During the Pre-Opening Period in the Paris Bourse," forthcoming in *Journal of Political Economy*.
- Bloomfield, R., and M. O'Hara, 1999, "Market Transparency: Who Wins and Who Loses?," *Review of Financial Studies*, 12, 5–35.
- Brooks, R. M., and T. Su, 1997, "A Simple Cost Reduction Strategy for Small Liquidity Traders: Trade at the Opening," *Journal of Financial and Quantitative Analysis*, 32, 525–540.
- Cao, C., E. Ghysels, and F. Hatheway, 2000, "Price Discovery without Trading: Evidence from the NASDAQ Pre-opening," forthcoming in *Journal of Finance*.
- Domowitz, I., and J. Wang, 1994, "Auctions as Algorithms: Computerized Trade Execution and Price Discovery," *Journal of Economic Dynamics and Control*, 18, 29–60.
- Dutta, P., and A. Madhavan, 1995, "Price Continuity Rules and Insider Trading," *Journal of Financial and Quantitative Analysis*, 30, 199–221.
- Economides, N., and R. A. Schwartz, 1995, "Electronic Call Market Trading," *Journal of Portfolio Management*, 21, 10–18.
- Edwards, A. K., 1999, "NYSE Specialists Competing with Limit Orders: A Source of Price Improvement," working paper, Securities and Exchange Commission.
- Forster, M., and T. George, 1996, "Pricing Effects and the NYSE Open and Close: Evidence from Internationally Cross-Listed Stocks," *Journal of Financial Intermediation*, 5, 95–126.
- Glosten, L. R., 1994, "Is the Electronic Open Limit Order Book Inevitable?," *Journal of Finance*, 49, 1127–1161.

- Garbade, K., and C. Sekaran, 1981, "Opening Prices on the New York Stock Exchange," *Journal of Banking and Finance*, 5, 345–355.
- George, T., and C.-Y. Hwang, 1997, "Information Flow and Pricing Errors: A Unified Approach to Estimation and Testing," working paper, University of Iowa.
- Hasbrouck, J., 1992, "Using the TORQ Database," working paper, New York Stock Exchange.
- Hasbrouck, J., and G. Sofianos, 1993, "The Trades of Market Makers: An Empirical Analysis of NYSE Specialists," *Journal of Finance*, 48, 1565–1594.
- Hasbrouck, J., G. Sofianos, and D. Sosebee, 1993, "New York Stock Exchange Systems and Trading Procedures," working paper, New York Stock Exchange.
- Ho, T., R. Schwartz, and D. Whitcomb, 1985, "The Trading Decision and Market Clearing under Transaction Price Uncertainty," *Journal of Finance*, 40, 21–42.
- Kavejecz, K., 1999, "A Specialist's Quoted Depth and the Limit Order Book," *Journal of Finance*, 54, 747–771.
- Kehr, C.-H., J. P. Krahnen, and E. Theissen, 1998, "The Anatomy of a Call Market: Evidence from Germany," working paper, Universität Frankfurt.
- Madhavan, A., 1992, "Trading Mechanisms in Securities Markets," *Journal of Finance*, 47, 607–642.
- Madhavan, A., and S. Smidt, 1993, "An Analysis of Changes In Specialist Quotes and Inventories," *Journal of Finance*, 48, 1595–1628.
- Madhavan, A., and G. Sofianos, 1997, "An Empirical Analysis of NYSE Specialist Trading," *Journal of Financial Economics*, 48, 189–210.
- Meier, R., 1998, "Benchmarking Analysis of Stock Exchange Trading, International Federation of Stock Exchanges (FIBV)," monograph, SBF-Bourse de Paris.
- New York Stock Exchange, 1991, *New York Stock Exchange Fact Book*, New York Stock Exchange, New York.
- Panchapagesan, V., 1999, "Identifying Specialist Trades in the TORQ Data—A Simple Algorithm," working paper, Washington University in St. Louis.
- Ronen, T., 1997, "Tests and Properties of Variance Ratios in Microstructure Studies," *Journal of Financial and Quantitative Analysis*, 32, 183–204.
- Schnitzlein, C. R., 1996, "Call and Continuous Trading Mechanisms Under Asymmetric Information: An Experimental Investigation," *Journal of Finance*, 51, 613–636.
- Schwartz, R. A., and R. A. Wood, 1995, "Dealer Markets, Derivative Expirations, and a Call," *Derivatives Quarterly*, 2, 38–45.
- Sofianos, G., 1995, "Specialist Gross Trading Revenues at the New York Stock Exchange," working paper, New York Stock Exchange.
- Sofianos, G., and I. Werner, 1997, "The Trades of NYSE Floor Brokers," Working Paper 97-04, New York Stock Exchange.
- Stoll, H., and R. Whaley, 1990, "Stock Market Structure and Volatility," *Review of Financial Studies*, 3, 37–71.
- U.S. Securities and Exchange Commission, 1963, *Report of the Special Study of Securities Markets, Part 2*, House Document 95, 88th Congress, 1st Session, Washington, D.C.
- Werner, I., 1998, "Who Gains From Steenths?," working paper, Ohio State University.
- Wohl, A., and S. Kandel, 1997, "Implications of an Index-Contingent Trading Mechanism," *Journal of Business*, 70, 471–488.
- Wyatt, E., 1999, "Nasdaq's Swings in Price Prompt Calls for Change," *New York Times*, July 8, 1999, C1.