PRICE EXPERIMENTATION AND SECURITY MARKET STRUCTURE

by

J. Chris Leach
Ananth N. Madhavan

20-92

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

The contents of this paper are the sole responsibility of the author(s).

Copyright © 1992 by J. Chris Leach and Ananth N. Madhavan.

Price Experimentation and Security Market Structure

J. Chris Leach
and
Ananth N. Madhavan*

The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

September, 1992

*We thank Lawrence Glosten, Mark Grinblatt, Michel Habib, Richard Kihlstrom, Paul Kupiec, Seymour Smidt, Chester Spatt (the editor), an anonymous referee, and the participants of workshops at Carnegie Mellon University, Cornell University, Pennsylvania State University, Wharton, and the European and American Finance Association meetings for their helpful comments. Assistance from the Geewax-Terker Research Fund is gratefully acknowledged. Any errors are entirely our own.

Abstract

This paper examines the role of market makers in facilitating price discovery. We show that a specialist may experiment with prices to induce more informative order flow, thereby expediting price discovery. Market makers in a multiple dealer system, unlike a specialist system, do not have the incentives to perform such costly experiments because of free-rider problems. Consequently, the specialist system may provide open markets where competition fails, but at the cost of wider bid-ask spreads. We analyze the effect of experimentation on the bid-ask spread and provide an exploratory analysis of intraday specialist data which is consistent with our price experimentation hypothesis.

1 Introduction

The recent organizational changes in European stock markets and the emergence of active financial markets in newly industrialized and former communist countries have focused attention on the importance of security market structure.[1] Most active securities trade in markets where market makers expedite trading by providing price quotations and supplying liquidity on demand. Integral to this function is *price discovery*, the process by which market-clearing prices are found. One would expect market makers, by virtue of their central position in the market, to play a critical role in price discovery. Further, differences in market microstructure are likely to have a substantial influence on the incentives for market makers to assist in price discovery. Yet, the function of price discovery has only recently been recognized in the literature.[2]

This discussion raises several fundamental questions: (1) Do market makers possess the ability to facilitate price discovery? (2) Do market makers have incentives to assist in price discovery if this activity is costly and the benefits are largely external? (3) Do differences in market structure affect the ability and incentives of market makers to aid price discovery? (4) How does price discovery affect the distribution of returns, the bid-ask spread, and transaction volume? (5) Are certain market structures preferred over others?

To focus attention on these aspects of market microstructure, we analyze and contrast the process of price discovery under two polar forms of market organization: a competitive multiple dealer system and a monopolistic specialist system. We analyze a particular strategy which a market maker can adopt to aid price discovery. The idea behind this strategy is simple. If trade originates partly from traders who possess private information and partly from traders with liquidity needs, then order flow conveys a noisy signal about the asset's future value. In the models of Glosten and Milgrom (1985), Kyle (1985), and Easley and O'Hara (1987), market makers update their beliefs after observing order flow and adjust prices accordingly. In these models, it is frequently assumed that liquidity or noise trade is exogenous. A more realistic assumption is that liquidity demands are price elastic. When order flow is endogenous, the precision of the signal created by that order flow is affected by the price quotations of the market maker. Intuitively, widening the bid-ask spread reduces the fraction of expected volume originating from noise traders, thereby making the observed order flow more informative. In contrast to previous microstructure models, market makers need not just learn passively from observed order flow but can strategically set quotes to induce the revelation of information.[3] In an intertemporal maximization problem, this link between signal quality and price quotes creates incentives for market makers to 'test the waters' by experimenting with prices in order to learn the full-information price more rapidly, even when such actions reduce immediate expected profits. We term such behavior *price experimentation.* Gammill (1990) makes a similar argument when he analyzes 'collusive' market structures which induce the perfect revelation of private information by providing bribes to insiders. Our model differs from Gammill's model in that market makers cannot distinguish between liquidity and informed traders, and consequently cannot perfectly infer the private information of insiders.

While market makers have the *ability* to expedite price discovery through such intertemporal pricing strategies, their *incentive* to undertake costly price discovery differs across various market structures. A specialist can set prices to induce costly informed trading, and recover his or her losses in later periods. Unlike specialists, competing dealers are unwilling to experiment at cost because the external nature of investment in the production of information creates a free-rider problem as other dealers observe the trading history. Indeed, it is possible that the specialist's monopolistic position may make future gains from trading with more precise knowledge sufficiently lucrative so that he chooses to keep the market open even though a multiple dealer system would close the market. This suggests a possible rationale for the continued existence of a specialist system. A related argument for the specialist system appears in Glosten (1989). Glosten considers a single-period model where order size is variable. If the level of noise trading is sufficiently low, Glosten shows that competitive dealers will be unable to make markets. However, Glosten shows that a specialist can make expected losses on large trades and expected profits on small trades with non-negative expected profits overall, thereby keeping the market open. Our model differs from Glosten in two important respects. First, while Glosten considers within-period subsidization from small to large trades, we consider a dynamic model with intertemporal subsidization from later traders to earlier traders. Second, unlike Glosten, the specialist may find it more profitable to exclude informed trade when competitive markets would permit such trade. A specialist system need not necessarily lead to more informative prices. From the viewpoint of public policy, there is a trade-off between the attributes of the two systems.

The model provides testable hypotheses concerning price dynamics, volatility, bid-ask spreads, and trading volume. We show that price experimentation is associated with a widening of the bid-ask spread, lower trading volume, and differences from the usual pattern of price formation. From an empirical viewpoint, however, testing our model is difficult because the implications of experimental behavior are consistent with other models of market maker behavior, and in particular, models of dynamic inventory control. In an attempt to address this ambiguity, we use intraday data obtained from a NYSE specialist. Our analysis suggests that experimentation may be empirically detectible even after corrections for inventory concerns and competition from the limit book.

The paper proceeds as follows: Section 2 introduces the theoretical framework. Section 3 provides a motivating example which is explicitly solved. Section 4 provides general results concerning the effect of price discovery on bid-ask spreads, volume, and price dynamics. In addition, we discuss the link between microstructure and market failure and the public policy implications of our analysis. Section 5 discusses implications of the price experimentation hypothesis and presents an exploratory empirical investigation of market maker experimentation. Section 6 concludes the paper with a summary. Proofs are in the Appendix.

2 The Theoretical Framework

Our model concerns the trading on short-lived information immediately preceding an information event such as a dividend or earnings announcement.[4] Formally, we consider the process of price discovery for a single risky security whose (stochastic) full-information price T periods in the future is denoted by $\tilde{\theta}$. One group of potential investors, informed traders, receive a private signal which reveals this full-information value. The remaining traders, referred to as liquidity traders, do not receive private information.

Trading takes place in T successive periods or trading rounds. All trading takes place with risk-neutral market makers, who stand ready to buy or sell the security on demand. In each period market makers set irrevocable bid and ask prices. Trade size is one round lot.[5] The market makers have imperfect information about the value of the asset, and learn from the pattern of order flow. After observing and participating in the trading activity, if any, for that round, the market maker updates his or her prior beliefs according to Bayes' Rule and sets quotes for the next trading round. After the trading activity, if any, in round T, the full-information value of the security is announced. To focus attention on the effects of market microstructure on the process of price formation, we will contrast two trading systems. In the first, referred to as the *specialist system*, there is one market maker who has a monopoly on market making in the risky security. By contrast, in the second system, referred to as the *multiple (competitive) dealer system*, there are at least two market makers who set prices competitively. The key distinction between these two systems lies in the fact that the specialist, unlike the competitive dealer, has the freedom to set prices that do not necessarily equate expected profits to zero in every period.

Let b_t and a_t be the bid and ask prices set for time $t = 1, \ldots, T$. The price pair selected at time t is denoted $p_t = (b_t, a_t)$, and these prices are chosen from an interval $P \subset \Re_+$ which contains the possible values of $\tilde{\theta}$. Let q_t denote the trading activity in round t where $q_t = -1$, $q_t = +1$ and $q_t = 0$ indicate a sell (to the market maker) at the bid, a purchase (from the market maker) at the ask, and no trading, respectively. Let the terminal value of the security be θ_1 with (prior) probability μ_1 and θ_2 with probability

$(1-\mu_1)$, where $\theta_1 < \theta_2$.

We still must describe how q_t is determined. Informed traders know with certainty the realization of $\tilde{\theta}$. Uninformed liquidity traders have reservation prices distributed with a cumulative distribution function $F(r)$, where then average, $\int r dF(r)$, is contained in the interval (θ_1, θ_2).[6] The reservation price of an uninformed trader is a simple way of representing the notion that these traders have other motivations for trading besides accurate information. Such motives include portfolio diversification, life-cycle consumption-savings decisions, and temporary liquidity shocks. They may also trade simply because they are unaware that an information event has taken place, in which case r would be interpreted as the expected value of the security given 'stale' information.

The composition of the trading population is such that the proportion λ of traders are informed and $(1-\lambda)$ are uninformed. Thus, the market maker is uncertain about the type of trader he or she faces unless the set prices exclude informed traders. For $b_t \leq \theta_1$ and $a_t \geq \theta_2$, only the uninformed will trade with the market maker. Informed and uninformed traders are selected at random from the trading population. We assume that both types of traders are drawn at random from a continuum of investors, so that the probability of being selected twice for trading is zero. This prevents strategic gaming of the market maker by traders.[7]

Upon observing (b_t, a_t), and the realization of $\tilde{\theta}$, denoted by θ, an informed trader executes trades according to:

$$q_t = \begin{cases} -1 \text{ (hit on bid)} & \text{if } b_t > \theta \\ +1 \text{ (hit on ask)} & \text{if } a_t < \theta \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

The price-sensitive uninformed trader, with a reservation price r, trades only if the trade is advantageous given his or her reservation price:

$$q_t = \begin{cases} -1 \text{ (hit on bid)} & \text{if } b_t > r \\ +1 \text{ (hit on ask)} & \text{if } a_t < r \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Since different uninformed traders may have different reservation prices in the underlying sub-population of uninformed, the expected trade execution given that an uninformed trader has been drawn is:

$$q_t = \begin{cases} -1 & \text{with probability } F(b_t) \\ +1 & \text{with probability } 1 - F(a_t) \\ 0 & \text{with probability } F(a_t) - F(b_t). \end{cases} \quad (3)$$

These best replies are known by the market maker who incorporates them into his maximization problem. Let χ_A be the indicator function on the set A, i.e., $\chi_A(x) = 1$ if $x \in A$ and 0 otherwise. Denote the history of trading at time t by $h_t = \{p_\tau, q_\tau\}_{\tau=1}^{\tau=t-1}$. The posterior at time $t+1$, μ_{t+1}, can be written using Bayes' Rule as:

$$\text{Prob}[\tilde{\theta} = \theta_1 | q_t = -1, b_t, a_t] \equiv \mu_{t+1}(-1) = \quad (4)$$

$$\frac{\mu_t[\lambda\chi_{\theta_1 < b_t} + (1-\lambda)F(b_t)]}{\mu_t[\lambda\chi_{\theta_1 < b_t} + (1-\lambda)F(b_t)] + (1-\mu_t)[\lambda\chi_{\theta_2 < b_t} + (1-\lambda)F(b_t)]}$$

and

$$\text{Prob}[\tilde{\theta} = \theta_1 | q_t = +1, b_t, a_t] \equiv \mu_{t+1}(+1) = \quad (5)$$

$$\frac{\mu_t[\lambda\chi_{\theta_1 > a_t} + (1-\lambda)(1-F(a_t))]}{\mu_t[\lambda\chi_{\theta_1 > a_t} + (1-\lambda)(1-F(a_t))] + (1-\mu_t)[\lambda\chi_{\theta_2 > a_t} + (1-\lambda)(1-F(a_t))]}.$$

Similarly, we obtain the posterior in the event of no trade:

$$\text{Prob}[\tilde{\theta} = \theta_1 | q_t = 0, b_t, a_t] \equiv \mu_{t+1}(0) = \quad (6)$$

$$\frac{\mu_t[\lambda\chi_{\theta_1 \in I} + (1-\lambda)(F(a_t) - F(b_t))]}{\mu_t[\lambda\chi_{\theta_1 \in I} + (1-\lambda)(F(a_t) - F(b_t))] + (1-\mu_t)[\lambda\chi_{\theta_2 \in I} + (1-\lambda)(F(a_t) - F(b_t))]},$$

where $I = [b_t, a_t]$.

The single-period reward (expected profit) function for the market maker in any period is given by:

$$\begin{aligned}\pi(b,a;\mu) = \mu \quad & \Big\{\lambda\left[\chi_{\theta_1<b}(\theta_1 - b) + \chi_{\theta_1>a}(a-\theta_1)\right] \\ & + (1-\lambda)\left[F(b)(\theta_1 - b) + (1-F(a))(a-\theta_1)\right]\Big\} \\ + (1-\mu) \quad & \Big\{\lambda\left[\chi_{\theta_2>a}(a-\theta_2) + \chi_{\theta_2<b}(\theta_2 - b)\right] \\ & + (1-\lambda)\left[F(b)(\theta_2 - b) + (1-F(a))(a-\theta_2)\right]\Big\}. \end{aligned} \tag{7}$$

With this framework in place, we can now discuss how prices are determined in the specialist and multiple dealer systems.

2.1 The Specialist System

The specialist is assumed to maximize his total trading profits. At first, it might appear that this problem, for a given prior distribution is static. However, this is not the case. The current placement and size of the bid-ask spread affect the distribution of order flow, and hence of next period's beliefs and rewards. In short, the data used as inputs into the learning process are endogenously determined. Consequently, a specialist who can set prices may choose to depart from maximizing current expected profits in order to produce information which is of greater future value than the immediate foregone profits. In doing so, the specialist experiments with prices by quoting a bid-ask spread to maximize the discounted sum of his profits.[8]

Formally, the market maker's maximization yields expected rewards denoted by:

$$V_n^*(\mu_1) = \sup_{\{b_t,a_t\}} E\left[\sum_{t=T-n+1}^{T} \pi(b_t, a_t; \mu_t)\right], \tag{8}$$

for n remaining trading rounds where the choice variables (a_t, b_t) are *policy functions* that are history-dependent, and the prior belief, μ_t, evolves according to Bayes' rule.[9] The expectation is taken over all random variables. The specialist recognizes that the selection of b_t and a_t has two effects: (1) the immediate impact on expected profits in round t and (2) the indirect effect through the posterior on expected rewards in the following rounds. This is demonstrated by the fact that in round $t+1$, b_t and a_t enter the problem explicitly through the functional form of μ_{t+1}. In Section 3, we explicitly solve a two-period problem which illustrates the basic intuition.

Note that by the Law of Iterated Expectations, the expected posterior is just the prior, i.e., $E[\mu_{t+1}|\mu_t] = \mu_t$. Even so, the specialist knows he or she will alter future actions as beliefs evolve and, therefore, chooses current period actions to extract information optimally from early rounds of trading. The location and spread of bid and ask prices dictates a level of investment in the production of information. Using the Bellman principle of optimality, we can write equation (8) as a functional equation:

$$V_T^*(\mu_1) = \sup_{\{a_1,b_1\}} \{\pi(a_1, b_1; \mu_1) + E[V_{T-1}^*(\mu_2(\tilde{q}))]\}. \tag{9}$$

The terminal condition is:

$$V_1^*(\mu_T) = \sup_{\{a_T,b_T\}} \{\pi(a_T, b_T; \mu_T)\}. \tag{10}$$

The function V_n^* represents the stochastic dynamic programming problem faced by the market maker who has n periods left to trade before $\tilde{\theta}$ is announced. The state variable for the problem is the market maker's prior *distribution*, the control variables are the bid and ask prices, and the transition equation is Bayes' rule. The following fundamental result about the value function allows us to simplify Equations (4), (5), (6) and (7) by ignoring quotes with $b_t > \theta_2$ or $a_t < \theta_1$ and will prove useful later:

Proposition 1 *For $\mu \in (0,1)$, the specialist's value function, $V_{T-1}^*(\mu)$, is convex and non-negative.*

The value function is non-negative because the specialist always has the option to quote prices at which no trade will occur, i.e., to close the market. Proposition 1 is important because the convexity of the value function implies that future information is valuable; the resolution of uncertainty in the prior distribution of μ makes the market maker better off. Since information is valuable, the market maker will never choose to set $b_t > \theta_2$ or $a_t < \theta_1$. Such quotes are strictly dominated by $b_t = \theta_2$ and $a_t = \theta_1$, respectively.

We now define an *optimal* or *experimental* price choice.

Definition 1 *An optimal (or experimental) first-period price choice* $p_1^* = (b_1^*, a_1^*) \in P$ *is the solution to the functional equation (9).*

This definition makes explicit the nature of price experimentation, i.e., a deliberate attempt by the market maker to set prices to produce information leading to more accurate beliefs about the true value of the traded security. We also refer to the optimal behavior of the specialist as *active learning.*[10] A market maker who recognizes the endogeneity of data used to revise his or her beliefs has incentives to facilitate price discovery. This happens not because the specialist cares about price accuracy *per se* but because more information helps make better future decisions, increasing his or her future profits.

In other microstructure models, market makers learn passively in the sense that they ignore the possibility of profiting from an increase in the precision of their estimate of the security's true value. Market makers condition their beliefs upon order flow using Bayes' rule but do not consciously *design* experiments to induce information revelation. Given their predominance in the microstructure literature, we use passive learning strategies as the benchmark for the specialist system.

Definition 2 *A first-period price choice,* $p_1^m = (b_1^m, a_1^m) \in P$ *is said to be non-experimental (or myopic) if* $p_1^m \in \arg\max \pi(b_1, a_1; \mu_1)$.

Such non-experimental prices would be used by a specialist who does not consider the strategic importance of endogenous information production in determining future trading profits.

2.2 The Multiple Dealer System

We turn now to the trading protocol in a competitive dealer market. As before, trading occurs sequentially, as in the specialist system, but the market maker with the lowest ask price or highest bid price gets to trade. Market makers observe all trades and quotes, whether or not they participate. Prices are determined through competition which minimizes expected profits subject to non-negativity. Before continuing, it is useful to note that the posterior functions for hits on the bid and ask sides are independent of the price chosen for the other side of the market. This provides separability of the two sides and using equation (7), we can write the expected profits from the ask side as:

$$\pi^a(a) = \lambda\chi_{\theta_2>a}(1-\mu)(a-\theta_2) + (1-\lambda)(1-F(a))(a - E_\mu[\tilde{\theta}]). \tag{11}$$

Similarly, we can write the expected profits from the bid side as:

$$\pi^b(b) = \lambda\chi_{b>\theta_1}\mu(\theta_1 - b) + (1-\lambda)F(b)(E_\mu[\tilde{\theta}] - b). \tag{12}$$

Definition 3 *Quote $p_1^c = (b_1^c, a_1^c) \in P$ is competitive if both of the following conditions are satisfied:*

$$\begin{aligned} \text{(a)} \quad b_1^c &= \sup\{b \in P : \pi^b(b) \geq 0\} \, . \\ \text{(b)} \quad a_1^c &= \inf\{a \in P : \pi^a(a) \geq 0\} \, . \end{aligned}$$

We use Definition 3 to characterize the equilibrium prices in the multiple dealer system:

Proposition 2 *In equilibrium, multiple dealer system quotes reflect only the benefit of learning on current trading profits. Specifically, the competitive bid and ask prices at time t are given by (respectively) $b_t^c = E[\tilde{\theta}|q_t = -1, b_t^c; \mu_t]$, $a_t^c = E[\tilde{\theta}|q_t = +1, a_t^c, \mu_t]$, where $a_t^c \geq b_t^c$.*

Proposition 2 is the Glosten and Milgrom (1985) result that the competitive bid and ask prices are *ex post* regret-free. Market makers in this system do not actively 'experiment' because they cannot profit in future trading and therefore are not concerned about the effect of their current actions on future trading. To design more informative experiments, dealers need to widen the quoted spread. In a multiple dealer system, a market maker who defects from (Betrand-Nash) equilibrium by raising the ask or lowering the bid will receive no trades. By contrast, the rational specialist takes into account the benefit of faster learning and increases the likelihood of current trade with the informed. This suggests that prices in the specialist system should converge more rapidly to their full-information value than competitive dealer prices. As we show later, this intuition is not necessarily correct.

3 A Motivating Example

To motivate the more general results that follow, we first digress to present a two-period example of the general model. There are two main reasons why this case is of interest. First, the optimization problem facing the decision maker can be characterized explicitly in the two-period case, helping to illustrate the nature of price experimentation. Second, the example permits us to investigate the conditions under which experimentation is likely to occur. Intuitively, we hypothesize that the deviation between experimental and non-experimental prices is likely to be greatest when (1) there is extreme uncertainty over asset values, as indicated by large $\theta_2 - \theta_1$, and (2) there is significant information asymmetry, represented by high values of λ. Our example demonstrates this intuition.

The two-period dynamic problem is:

$$\begin{aligned}\max_{\{a_1,b_1\}}\{\pi(b_1,a_1;\mu_1) &+ \gamma_1 \max_{\{b_2,a_2\}}(\pi(b_2,a_2;\mu_2|q_1=-1,\ldots)) \\ &+ \gamma_2 \max_{\{b_2,a_2\}}(\pi(b_2,a_2;\mu_2|q_1=+1,\ldots)) \\ &+ \gamma_3 \max_{\{b_2,a_2\}}(\pi(b_2,a_2;\mu_2|q_1=0,\ldots))\}\end{aligned} \tag{13}$$

where:

$$\begin{aligned}\gamma_1 &= \lambda(\mu_1\chi_{\theta_1<b_1} + (1-\mu_1)\chi_{\theta_2<b_1}) + (1-\lambda)F(b_1) \\ \gamma_2 &= \lambda(\mu_1\chi_{\theta_1>a_1} + (1-\mu_1)\chi_{\theta_2>a_1}) + (1-\lambda)(1-F(a_1)) \\ \gamma_3 &= \lambda(\mu_1\chi_{b_1\le\theta_1\le a_1} + (1-\mu_1)\chi_{b_1\le\theta_2\le a_1}) + (1-\lambda)(F(a_1)-F(b))\ .\end{aligned}$$

This equation states that the maximum is the first-period profit from the price choice (b_1, a_1) plus the expectation over the three possible posteriors of the second period's expected profits. By considering F uniform on the interval $[\theta_1, \theta_2]$ the first-order conditions become cubic polynomials in a_1 and b_1. For a base-case with $\theta_1 = 1$, $\theta_2 = 2$, $\lambda = 0.05$ and $\mu_1 = 0.5$, the market is open under all three regimes and $a^* \geq a^m \geq a^c$. The optimal ask price is \$1.764, which is 15.6% higher than the competitive ask quote and 0.04% higher than the non-experimental ask quote. The small difference between non-experimental and optimal quotes is a reflection of the two-period assumption; the market maker has just one trade to recover his investment in information production. In a more realistic setting, the number of trading opportunities prior to an information event may be quite large, increasing the future rewards to experimentation without affecting its cost.[11] Additionally, one might be pessimistic about the prospects of detecting price experimentation in actively traded securities when the ratio of the precision of public to private information is high. However, as uncertainty regarding fundamental values increases, or the level of liquidity-motivated trading declines, price experimentation plays an increasingly important role in the optimizing specialist's price quotes and therefore transaction prices.

Figure 1 demonstrates this intuition for our example. Active experimentation (measured by the ratio of the optimal ask price to the non-experimental ask price) becomes more marked as adverse selection worsens, i.e., as the probability of trading with an informed trader (λ) increases. All three types of ask prices rise as the adverse selection problem worsens, and, since competitive prices rise more quickly than specialist prices, the gap between competitive and specialist quotes decreases as adverse selection worsens. Figure 2 shows that departure from both non-experimental and competitive prices is more pronounced when uncertainty, as measured by mean-preserving increases in the distribution of θ, increases.

Our example also demonstrates that price experimentation is associated with systematic deviations from 'normal' price patterns. Here, the non-experimental and competitive quotes are known linear functions of the prior expected value of the asset and therefore will follow martingales. It is straightforward to show that the experimental quotes do not follow a martingale, provided the specialist system is open.

To establish this claim, suppose the specialist experiments, so that $a_1^* \neq a_1^m$. In the final period, the non-experimental price functional form and the optimal price functional form coincide and are therefore both linear in beliefs, μ_T^m and μ_T^*, respectively. Then, $E[a_T^*(\mu_T^*)] = a_T^*(E[\mu_T^*])$ and $E[a_T^m(\mu_T^m)] = a_T^m(E[\mu_T^m])$. By the martingale property of beliefs $E[\mu_T^*] = E[\mu_T^m] = \mu_1$, which implies $a_T^*(E[\mu_T^*]) = a_T^m(E[\mu_T^m])$ By the previous two equalities, $E[a_T^*(\mu_T^*)] = E[a_T^m(\mu_T^m)]$. For optimal prices also to follow a martingale, it must be that $a_1^* = a_1^m$ which contradicts our supposition that the specialist is experimenting. The example shows that price experimentation (and not just the exercise of market power) generates price paths which systematically depart from the martingale property of prices implicit in many microstructure models.[12] From an empirical perspective, this result is important because there is evidence that intraday prices do not follow martingales (see, e.g., Hasbrouck (1991)). The only alternative explanation for this finding is the inventory control hypothesis, for which the empirical evidence appears to be mixed. (See, e.g., Madhavan and Smidt (1991) and the references contained therein.) We now return to our analysis of the model under general assumptions regarding distributions and horizons.

4 Market Microstructure and Price Discovery

4.1 A Comparison of Multiple-Dealer and Specialist Systems

With our analytical structure in place, we can compare the specialist and dealer systems for general finite horizons and parametric assumptions. Our focus will be on the bid-ask spread, trade volume, and the information content of trades. We assume that F does not change over time.[13] We begin by considering general F's, but will introduce assumptions about the support of F as it becomes necessary to sharpen some of our results. First, we need to define market failure:

Definition 4 *At time t, the market is open (closed) on the ask-side if the ask price allows (excludes) profitable informed trade: $a_t < \theta_2$ ($a_t \geq \theta_2$). Similarly, the market is open (closed) on the bid-side if the bid price allows (excludes) profitable informed trade: $\theta_1 < b_t$ ($b_t \leq \theta_1$). The market is open if it is open on at least one side.*

From this point forward, we will suppress the subscript for $t = 1$. Our intuition (and the previous example) suggests that a specialist exploits his monopolistic position by setting wider spreads than those set by competitive dealers. This is generally true for our model when specialists and competitors open the market.

Proposition 3 *If the multiple dealer and specialist systems are open, then both the non-experimental and optimal bid-ask spread in a specialist system are wider than the competitive bid-ask spread. Specifically:* $b^* \leq b^c$, $b^m \leq b^c$, $a^c \leq a^*$ *and* $a^c \leq a^m$. *When* a^m *and* b^m *are unique,* $b^* \leq b^m \leq b^c$ *and* $a^c \leq a^m \leq a^*$.

Corollary 1 *Given beliefs* $\mu_t \in (0,1)$, *expected transaction volume* ($E[|q_t|]$) *is lower in a specialist system than in a multiple dealer system if the two markets are open. When* a_t^m *and* b_t^m *are unique, volume in the specialist system is lower under experimentation than under myopia.*

The intuition behind the proposition is as follows. A non-experimenting specialist will not quote a spread within the competitive bid-ask spread because, by the definition of b^c and a^c, this action produces negative expected profits. A dynamically optimizing specialist will always seek a more informative experiment than a specialist who does not experiment. If the probability of an informed trader's placing an order is positive, i.e., the market is open, the specialist will decrease the probability of trade with the uninformed through a wider (relative to the non-experimental and competitive) spread.[14] Note that, by definition, expected current profits are lower for the optimal bid and ask prices than for the non-experimental ones. The specialist's market power implies that spreads exceed those set by competitive dealers, but the dynamically optimizing specialist continues to widen the spread beyond the profit-maximizing myopic spread to create an additional investment in the production of information.

The corollary follows directly from Proposition 3 and the structure of uniformed trading. Importantly, experimentation is characterized by lower expected initial volume. While it might seem plausible that higher volume gives rise to a more accurate signal, in our model, the increase in volume accompanying a decrease in the spread is related only to uninformed trade and acts to diminish the quality of the signal.

4.2 Market Failure

We turn now to a comparison of the robustness of the two systems under severe adverse selection. In particular, we are interested in whether a specialist system can function when a multiple dealer system does not, and vice versa. We separate our results by our assumptions on F, the distribution of uninformed reservation prices.

4.2.1 F's Concentrated in $[\theta_1, \theta_2]$

Proposition 4 *For F concentrated in $[\theta_1, \theta_2]$, if the market is open under a multiple dealer system, it is optimal to open it under a specialist system. A necessary and sufficient condition for an open multiple dealer market is that first period expected profits are nonnegative for some bid-ask pair (not necessarily competitive) where at least one of the quotes is contained in (θ_1, θ_2).*

This result generalizes the price behavior demonstrated in our example. The specialist system is at least as robust as the dealer system for these F's. Even though by Proposition 3 we know that the specialist quotes a wider spread, it is not sufficiently wide to exclude informative trade which would be permitted under a multiple dealer system. This proposition does not discuss the behavior of a specialist under the situations where the competitive system is closed. In other words, we have shown that the specialist system is at least as robust, but is it more robust? For *F's* concentrated in $[\theta_1, \theta_2]$ the answer is 'yes.'

A simple example demonstrates this claim. Let $\lambda = \mu = 0.5$, $\theta_1 = 0$, $\theta_2 = 1$, and suppose that liquidity trader reservation prices are massed at $\frac{3}{8}$ and $\frac{5}{8}$, each with probability $\frac{1}{2}$. Consider the bid side of the market as the ask side is symmetric. We have:

$$\pi^b(b) = \frac{1}{4}\chi_{b>0}(-b) + \frac{1}{2}\left[\frac{1}{2}\chi_{\frac{3}{8}\leq b} + \frac{1}{2}\chi_{\frac{5}{8}\leq b}\right]\left(\frac{1}{2} - b\right).$$

Note for $b \geq \frac{1}{2}$, the market maker expects to lose to both informed and uninformed traders. For $b < \frac{3}{8}$, the market maker trades only with the informed. For competitors, we can therefore restrict attention to $\frac{3}{8} \leq b \leq \frac{1}{2}$. For this range of b we have:

$$\pi^b(b) = -\frac{1}{2}b + \frac{1}{8} < 0.$$

which demonstrates that a multiple dealer system closes. We must show that it will be open under a monopolistic specialist. Consider setting prices at $b = \frac{1}{4}$ and $a = \frac{3}{4}$. The single period expected loss from trading with an informed trader is $\frac{3}{8}$, but induces perfect revelation. Per period profits under perfect information are strictly positive as can be seen by considering $b = 0$, $0 < a < \frac{5}{8}$ for $\theta = 0$ and $a = 1$, $\frac{3}{8} < b < 1$ for $\theta = 1$. Define $0 < \bar{\pi} \leq 1$ to be the maximum per period profit which, by symmetry, is the same for $\theta = 0$ and $\theta = 1$. Then, the total expected profits with $n-1$ periods remaining is $(n-1)\bar{\pi}$ giving a value of the program under these prices of:

$$\lambda\left(-\frac{3}{8} + (n-1)\bar{\pi}\right) + (1-\lambda)V^*_{n-1}(\mu).$$

Since $V^*_{n-1}(\mu) \geq 0$, due to the option of closing the markets, this value will be positive (and the market is open under a specialist system) if $(n-1) > \frac{3}{8\bar{\pi}} > 0$. Letting n equal the smallest integer greater than $\frac{3}{8\bar{\pi}}$ completes the construction.

Our demonstration shows that a trade-off across the two structures does arise. A specialist system may be viable in a market where the adverse selection problem is so severe that a multiple dealer system cannot operate without suffering expected losses. Unlike a specialist system, competing market makers face a severe free rider problem — they have no incentives to experiment at cost because of the external nature of investment in the production of information. From Proposition 2, a dealer cannot recover the investment in information made by opening the market while quoting prices that produce an expected first-period loss because any future benefits are eliminated by competition as all dealers observe the trading history. Similarly, a specialist who does not experiment will not open a market which does not offer current profits. However, a dynamically optimizing specialist may incur the loss in return for increases in future trading profits with the uninformed. Previously, the more informative experiment designed by the optimizing specialist involved widening the bid-ask spread relative to the non-experimental and competitive cases (see Proposition 3) when that market is open. Following in the same line of reasoning, when it is closed, the only (more) informative experiments available involve the possibility of trading with the informed through a narrowing of the bid-ask spread. Although this increases the incidence of liquidity trading, order flow conveys signals that are absent when the market is closed. In summary, a specialist who experiments may permit informed trade when competitors and a non-experimenting specialist would close the market.

4.2.2 F's Which are not Concentrated in $[\theta_1, \theta_2]$.

We have found a class of F's for which the specialist system demonstrates superior robustness. Next we address the question of whether F's outside of this class permit cases where the specialist closes a market which would open under competition. This would arise when the specialist is content trading only with liquidity traders whose reservation prices are no longer restricted to the interval $[\theta_1, \theta_2]$. An optimizing specialist may choose to set bid and ask prices which preclude informed trading even though profits are positive under some quotes which open the market.

Proposition 5 *The specialist market (under both experimental and non-experimental behavior) may be closed even if the market is open under a multiple dealer system; a necessary condition for this event is that F is not concentrated in $[\theta_1, \theta_2]$.*

This proposition illustrates the difference between our model and Glosten (1989). Here, the distortion created by the market power of market makers may, in some situations, offset the gains from their ability to engage in intertemporal price discovery. The extent to which this trade-off is significant is an empirical question. Our results suggest that losses due to higher spreads may be offset by the greater robustness of the specialist system to adverse selection problems. The dominating effect depends on the structure of liquidity-motivated trading. Note that if the non-experimental quotes contain the competitive quotes and both markets are open, price discovery is still faster with a specialist system, even if the specialist does not experiment.

Although we have motivated our analysis within a short-horizon framework, the intuition of our model generalizes to longer horizons. A longer horizon increases the opportunities to use the information generated through experimentation. Even if insiders have retrading opportunities and a sufficiently long horizon to introduce the possibility of manipulating the specialist's beliefs, trade will still be informative and the bid-ask spread will still affect the level of informativeness, albeit in a more complex fashion than we have described. It appears unlikely that the (equilibrium) solution to the market maker's dynamic optimization problem with endogenous data would coincide with the myopic choice.

4.2.3 Price Discovery and Public Policy

Turning to the policy implications of our analysis, we note that Propositions 3 and 4 suggest that there may not be a form of market organization that provides both low bid-ask spreads and robustness to extreme problems of asymmetric information. The multiple dealer system is less robust to problems of asymmetric information than a specialist system; the specialist may choose to keep markets open in order to facilitate price discovery in times of extreme uncertainty. This suggests that a specialist system is less likely to fail during a period of extreme information asymmetry than a dealer market, although certainly there are other factors that influence these decisions.

However, our results allow the possibility of the opposite case for particular parameter values of the model. This ambiguity in our stylized model makes it impossible to make definitive normative statements regarding market structure. While price discovery has aspects of a socially beneficial externality, it is associated with costs in the form of wider bid-ask spreads. These results suggest the observed diversity of trading structures may reflect adaptations to different market conditions. Judging the relative costs and benefits of price discovery is not possible without complete knowledge of the parameters of the economy.

Our model does, however, yield some specific recommendations for policy designed to reduce volatility. Following the crash of 1987, several independent investigative commissions recommended the use of coordinated trading halts or circuit breakers to halt trading in the face of large price movements. Our analysis shows that these suggestions may be misguided. We show that 'abnormal' price movements may reflect price discovery by the market maker which helps the market to resolve uncertainty more rapidly. These movements are associated with a widening of the spread and a reduction in trading volume. Even when market failure appears imminent due to problems of asymmetric information or insufficient liquidity, closing the market simply makes it more difficult to re-open at a future date since there is no method for market makers to learn about the full-information price by observing order flow.[15]

5 The Empirical Evidence

5.1 Prices, Bid-Ask Spreads, and Volume

Our theoretical analysis leads to two natural questions: (1) How can we empirically detect price discovery through price experimentation? (2) Is there any evidence that price experimentation affects bid-ask spreads, transaction volume, and the distribution of returns in practice? The model developed in this paper relies upon a feedback effect from dealer price quotations to order flow. Thus, a logical starting point for an empirical investigation is to verify that there is indeed such a relation. Hasbrouck (1991), estimating a vector autoregressive model using intraday data for NYSE stocks, finds that revisions in quote midpoints Granger-Sims cause trades. Hasbrouck finds a negative relation between trades and lagged quote revisions and concludes that this finding is consistent with both inventory control models and price experimentation. Evidence of a similar nature is provided by Madhavan and Smidt (1992) who, using daily data for 16 NYSE stocks, find that quote revisions are related not only to subsequent trades, but also to changes in specialist inventory positions. While these results provide support for the basic feedback effect from quote revisions to market maker inventory which underlies our theoretical model, they are far from conclusive because of the problem of observational equivalence with the inventory control hypothesis.[16] Further, these studies provide only indirect support for our hypotheses as they examine only changes in quote midpoints, and not in the size of the spread.[17]

A more direct approach to determining whether there is empirical evidence of price experimentation is to examine a time-series of intraday quote revisions and trades. Our model predicts that experimentation is associated with an abnormal widening of the bid-ask spread and a reduction in trading volume. Conjecturing that price experimentation may directly follow a rapid change in price triggered either by public announcements or by order imbalances, Handa (1991) uses intraday transaction data for approximately 2,000 NYSE and AMEX stocks for 1988 and 1989 to examine the future distribution of price changes conditional upon the history of quote revisions. He finds that as the placement of bid-ask prices changes, the size of the spread doubles from \$0.203 to about \$0.410, which is consistent with experimentation triggered by an abnormal price movement. Further, the conditional distribution of prices violates the martingale property conditional upon a price movement (a finding consistent with the arguments in our example section). Handa also finds that conditional on no change in the mid-point of the spread, the bid price rises by 0.27 cents and the ask price falls by 0.31 cents. This finding is difficult to explain using either the information or inventory models, but can be explained by our model. Consider the two-period example of Section 3. In the absence of trade in the first period, beliefs are unchanged and consequently there is no change in the non-experimental prices in the final period. However, the experimental prices in the second period coincide with the non-experimental prices, implying a narrowing in the size of the bid-ask spread without any movement in the placement of the mid-quote.

While this evidence is certainly suggestive, there could be other explanations for these findings. In particular, as noted by Hasbrouck (1991), models of inventory control give rise to similar predictions regarding intraday variation in spreads. For example, in the model of Amihud and Mendelson (1980), larger dealer holdings imply a widening in the quoted spread together with associated price movements. Similarly, intertemporal variation in competition to the specialist from floor traders or the limit order book may explain changes in the bid-ask spread. We turn now to an empirical examination of the experimentation hypothesis advanced in this paper which attempts to control for changes in spreads attributable to factors other than experimentation. The study is exploratory in nature; it is not intended to be a definitive test of the experimentation hypothesis.

5.2 An Empirical Examination of Price Discovery and Spreads

The data used here consist of almost 75,000 transactions of a New York Stock Exchange (NYSE) specialist firm from February 1, 1987 to December 31, 1987 in 16 stocks. These data were combined with data from the Institute for the Study of Securities Markets (ISSM) on bids, asks, transaction prices, and trading volume, to provide a complete time-series of trading activity in the specialist's stocks across all domestic markets. Madhavan and Smidt (1991) use these data to analyze specialist trading activity, and provide a complete description of the data and the procedures used to verify their accuracy.

The data are ideally suited for an empirical study of price discovery for two reasons. First, the NYSE operates a specialist system with a single market maker for each stock, and thus we are more likely to detect short-run experimentation with these data. Second, the data are sufficiently detailed that we are able to test the model's major prediction that market makers strategically widen the bid-ask spread to enhance price discovery, while we control for the influence of factors other than experimentation. In particular, the specialist data contain information on inventory positions during the day. Further, the ISSM data also contain information on the depth of the bid and ask quotes, i.e., the number of shares that can be accommodated at the prevailing quotes. The reported depth reflects the liquidity provided by both the specialist and limit orders at the prevailing quotes. Market depth thus proxies for the thickness of the limit order book, and allows us to control for changing competitive pressures.

Our focus here (unlike Madhavan and Smidt (1991) who examine transaction prices) is on the determinants of price quotations. From an empirical viewpoint, the use of quotations creates some difficulties because of the large numbers of duplicate quotation records; typically, quotations are recorded at regular intervals, even if there are no changes in the bid or ask prices.[18] Accordingly, we classified the data into 6 hourly periods from the opening of the NYSE at 9:30 am to 3:30 pm, EST, and one half-hour period from 3:30 pm to the closing at 4:00 pm, EST. For period t, we record the first quoted bid and ask prices (denoted by a_t and b_t, respectively), the average of the bid and ask depth associated with the first quotes (denoted by D_t) of the period, and the specialist's inventory position at the start of the hour (denoted by I_t). It is important to note that the first quoted bid and ask prices are the so-called Best Bid and Offer (BBO) quotes, i.e., the highest bid and smallest ask prices available in the National Market system. Thus, these quotes may reflect prices set by limit orders. To the extent that this is the case, the results of the empirical analysis are biased against detection of active price discovery because this limits the ability of the specialist to experiment by enlarging the spread.

Denote by $p_t = \frac{a_t+b_t}{2}$ the average of the bid and ask quotes in hour t, and let $s_t = \frac{a_t-b_t}{p_t}$ denote the corresponding percentage bid-ask spread. Finally, let $r_t = |p_t - p_{t-1}|$ represent the absolute revision in the specialist's quoted prices from period $t-1$ to period t.

The specialist's incentives to engage in price discovery are greatest when price volatility is large and beliefs about the security's fundamental value are imprecise. Thus, other things being equal, experimentation is most likely to occur in the early part of the trading day and following large absolute price movements.

Thus, our model would predict a widening of the bid-ask spread, s_t, following an increase in price volatility, measured by the absolute return, $|r_t|$.[19] Further, because experimentation is likely to occur after the overnight non-trading period, we expect to find that, other things being equal, the bid-ask spread, s_t, is larger in the opening hour.[20] In turn, the model implies that the specialist's beliefs, and hence his price quotations, are more likely to be revised following a period of price experimentation. Thus, we would expect an increase in the quoted bid-ask spread, s_t, to precede an increase in $|r_{t+1}|$. Note that r_t is measured as the revision in quotes, not the change in transaction prices, so as to eliminate potential biases introduced by bid-ask bounce. As noted above, inventory effects may also account for this relation. The inventory control hypothesis implies that the quoted spread, s_t, increases with larger (absolute) inventory levels; similarly, greater competition from limit orders should narrow spreads, implying that s_t is inversely related to the average market depth, D_t. Finally, as noted above, there may be important time-of-day effects on bid-ask spreads. Let O_t and C_t represent dummy variables for the opening and closing intervals, respectively. The discussion above suggests that experimentation is most likely to occur early in the day and is least likely at the end of the day, so that we expect s_t to be positively related to O_t and negatively related to C_t.

This motivates the following regression:

$$s_t = \beta_0 + \beta_1|r_{t+1}| + \beta_2|r_t| + \beta_3|I_t| + \beta_4 D_t + \beta_5 O_t + \beta_6 C_t + \epsilon_t \tag{14}$$

where ϵ_t is an error term. Based on the discussion above, we expect that $\beta_0, \beta_1, \beta_2, \beta_3$, and β_5 are positive while β_4 and β_6 are negative.

Table 1 presents the results from estimating equation (14) for the 16 stocks (ordered from lowest to highest by transaction frequency) using the method of unconditional least-squares (the Yule-Walker procedure) to correct for first-order serial correlation.[21]

In Table 1, the constant, β_0, represents the time-invariant portion of the spread, which is attributable to transaction costs; it varies between 0.56% and 2.6%, and is highly significant. Our model predicts that, controlling for inventory effects and competitive pressures, the coefficient of $|r_{t+1}|$ is positive, i.e., that a widening of the spread is a predictor of quote revisions over the next hour. Indeed, β_1 is positive in all but two cases and is significant in 13 cases. Additionally, as hypothesized, the coefficient β_2 is positive and significant in 11 cases. While such widening of the spread following a large quote revision is predicted by our model, it is also consistent with inventory control and asymmetric information theories.

The coefficient of absolute inventory, $|I_t|$, is positive for 15 stocks and is statistically significant for all but one of these stocks. There are strong inventory effects on the size of bid-ask spreads. Further, there is strong evidence that spreads are inversely related to market depth, which is consistent with the hypothesis that competition from limit orders reduces quoted bid-ask spreads. In all but one case, the coefficient β_4 is negative, and in 10 of these cases it is significant. The time-of-day dummy variables suggest that spreads are widest at the opening and narrow at the end of the day. As the opening follows a long period of non-trading, the experimentation hypothesis would predict that spreads are widest at that time. This result could also reflect other factors such as higher perceived information asymmetries or risk at the opening.[22] The effect is strong; β_5 is positive for all stocks and is significant for 10 stocks. The coefficient β_6 is negative in all but one case, but is statistically significant in only one case. This suggests that the increase in the end-of-day quotes documented by some researchers (see, e.g., Harris (1986) and Harris (1989)) may reflect factors such as inventory that we control for in our analysis of intraday spreads. Finally, the model does quite well at explaining the temporal variation in spreads; the regression R^2 is high, 49% on average, with a minimum of 29% and a maximum of 71%.

To summarize, our examination of intraday variation in bid-ask spreads provides evidence consistent with the model, even after correcting for other factors that affect spreads. This study is not intended to be a definitive test of the theoretical model; a complete empirical investigation is beyond the scope of this paper. However, there are several logical avenues for extending the analysis of this section. One approach would be to examine quotations and volume at times when experimentation is likely to occur, such as before and after known information events (such as dividend announcements) or conditional upon an increase in volatility. For stocks, such shifts in volatility could be detected using options data or through an analysis of errors from a pricing model. Another approach would be to perform tests across different securities in the same market or across different markets. Within the NYSE, for example, there are significant differences in the level of competition faced by specialists. For very active stocks, competitive pressures may preclude any active price discovery by the specialist, whereas for thinly traded stocks where the specialist faces little competition, such behavior may have a significant effect on prices. Thus, our model would suggest that the unconditional time-series properties of quotes and volume is systematically related to the specialist's participation rate.

It may also be possible to test the model by examining securities traded in different markets. Both common stocks and listed stock options are traded in markets where both specialist and multiple dealer systems are in use.[23] Since experimentation is more likely to be observed in a specialist system than a multiple dealer system, one can test whether the model's time-series predictions regarding quotations and volume differ systematically across these market structures. These are topics for further research.

6 Conclusions

This paper analyzes the process of intertemporal price formation under two forms of market organization: a multiple dealer system and a specialist system. The focus is on the role of market makers in *price discovery*, i.e., the process of finding the full-information price of the security. The crucial difference between the two systems stems from the fact that, unlike a multiple dealer market, prices in a specialist system need not be set to equate expected profits to zero on every trade. This feature affects price discovery in a fundamental way.

A rational market maker recognizes that the data used in decision making in the future are endogenously determined because current actions affect the probability distribution of order flow. Consequently, a market maker may undertake myopically suboptimal behavior in order to induce more informative order flow, thereby expediting price discovery. This type of *price experimentation* constitutes investment in the production of information. Although experimentation is costly in the short-run, it can be optimal in a dynamic context because more information leads to more profitable future actions. By contrast, competitive dealers do not have the discretion to make such investments because of the zero expected profit condition. The external nature of investment in the production of information creates a free-rider problem when there is competition among market makers. With perfect floor information, no single competitive dealer has the incentive to experiment at cost in the early trading rounds as there is no opportunity to recover the investment in information production in later trading.

Our model yields a number of testable hypotheses regarding price formation. Recent studies using intraday transaction data provide evidence in favor of the price experimentation hypothesis. We provide a new empirical investigation which shows that these effects are still present when we adjust for inventory considerations. From the viewpoint of public policy, our results suggest a subtle trade-off between the two systems. The specialist's ability to discover prices through experimentation can provide the impetus for open markets where multiple dealer systems fail. In general, however, the specialist's bid-ask spread will be wider than the competitive dealer's spread. Further, a specialist may find it optimal to set a spread that precludes all learning while learning opportunities exist under the multiple dealer system. These results suggest the observed diversity of trading structures may reflect adaptations to different market conditions.

Appendix

Proof of Proposition 1:

Non-negativity follows from the fact that the market maker always has the option of quoting prices outside the interval $[\theta_1, \theta_2]$ now and in every future period. (For F's with support outside of $[\theta_1, \theta_2]$, there will be a positive lower bound.) This establishes zero as a lower bound on the value function for all μ. To establish convexity, consider a $\mu' \in [\mu, \mu'']$ for $\mu, \mu'' \in [0, 1]$ and $\mu'' \geq \mu$. Note that:

$$\mu' = \phi\mu + (1-\phi)\mu'' \tag{A.1}$$

where:

$$\phi = \frac{(\mu'' - \mu')}{(\mu'' - \mu)}. \tag{A.2}$$

Suppose there were an action or 'experiment' which would induce a posterior of μ with probability ϕ and μ'' with probability $1-\phi$ when the prior is μ'. This experiment is clearly more informative than (sufficient for) the alternative of no experiment since it may lead to a revision in beliefs. Since one cannot be worse off on average by learning the outcome of the experiment, it must be that:

$$\phi V^*_{T-1}(\mu) + (1-\phi)V^*_{T-1}(\mu'') \geq V^*_{T-1}(\mu'), \tag{A.3}$$

which establishes convexity. ∎

Proof of Proposition 2:

Using Bayes' rule we obtain:

$$E[\tilde{\theta}|q = +1; a] = \frac{(1-\lambda)(1-F(a))E_\mu[\tilde{\theta}] + \lambda\chi_{\theta_2 > a}(1-\mu)\theta_2}{\lambda\chi_{\theta_2>a}(1-\mu) + (1-\lambda)(1-F(a))}. \tag{A.4}$$

Solving for $E_\mu[\tilde{\theta}](1-\lambda)(1-F(a))$, substituting it into equation (11) and rearranging, we obtain:

$$\pi^a(a) = (a - E[\tilde{\theta}|q = +1; a])[\lambda\chi_{\theta_2>a}(1-\mu) + (1-\lambda)(1-F(a))]. \tag{A.5}$$

Since $\text{Prob}[q = +1] = \lambda\chi_{\theta_2>a}(1-\mu) + (1-\lambda)(1-F(a))$, equation (A.5) can be written as:

$$\pi^a(a) = \text{Prob}[q = +1](a - E[\tilde{\theta}|q = +1; a]).$$

Similarly we obtain:

$$\pi^b(b) = \text{Prob}[q=-1](E[\tilde{\theta}|q=-1;b]-b).$$

and the result follows immediately.∎

Proof of Proposition 3:

First, we note that $\mu_2(-1)$ and $\mu_2(+1)$ are decreasing in b and a respectively for $\mu \in (0,1)$ when the markets permit informed trade. Fix a b and suppose $a^c > a^*$. Then $\mu_2(+1,a^*) \geq \mu_2(+1,a^c)$ and $\mu_2(-1)$ is the same for both choices of a. From Proposition 1 (or Blackwell's Lemma), this fact implies that $\mathrm{E}\left[V^*_{T-1}(\mu_2(b,a^c,\tilde{q}(a^c)))\right] \geq \mathrm{E}\left[V^*_{T-1}(\mu_2(b,a^*,\tilde{q}(a^*)))\right]$ for any b. To avoid a contradiction of the definition of a^* we must have $0 \leq \pi(b,a^c,\mu) \leq \pi(b,a^*,\mu)$. But this shows that a^* is a price lower than a^m which has non-negative profits - contradicting the property of a^c as the lowest ask price yielding non-negative expected profits from the ask side. Therefore, it must be that $a^* \geq a^c$. Next suppose $a^c > a^m$. By the definition of a^c expected profits at a^m must be negative and this contradicts the definition of a^m.

Suppose now a^m is the unique non-experimental maximizer and assume in order to produce a contradiction that $a^m > a^*$. We have $\pi(b,a^m,\mu) > \pi(b,a^*,\mu)$ by definition of a^m and $\mathrm{E}\left[V^*_{T-1}(\mu_2(b,a^m,\tilde{q}(a^m)))\right] \geq \mathrm{E}\left[V^*_{T-1}(\mu_2(b,a^*,\tilde{q}(a^*)))\right]$ by our supposition. But this contradicts the definition of a^*. We conclude that $a^c \leq a^m \leq a^*$.

The bid side is analogous and the spread implications follow immediately.∎

Proof of Proposition 4: By hypothesis, at least one of the competitive prices is in (θ_1,θ_2). Excluding informative trade by pricing $b \leq \theta_1$ and $a \geq \theta_2$ produces a degenerate distribution in posteriors at $\mu_2 = \mu$. The distribution of posteriors arising from the competitive prices are a mean-preserving spread of this distribution. By the convexity of V in μ (Proposition 1) we know $\mathrm{E}[V^*_{T-1}(\mu_2(b^c,a^c,\tilde{q}(b^c,a^c)))] \geq \mathrm{E}[V^*_{T-1}(\mu_2(\theta_1,\theta_2,\tilde{q}(\theta_1,\theta_2)))] = \mathrm{E}[V^*_{T-1}(\mu)]$. We also know that $\mathrm{E}[V^*_{T-1}] \geq 0$. Since competitive prices yield zero rewards under our assumptions on P, we know $\pi(b^c,a^c,\mu) = 0$. Therefore, b^c and a^c are prices which open the market and we have:

$$[\pi(b^c,a^c,\mu) + \mathrm{E}[V^*_{T-1}(\mu_2(b^c,a^c,\tilde{q}(b^c,a^c)))] \geq \pi(\theta_1,\theta_2,\mu) + \mathrm{E}[V^*_{T-1}(\mu_2(\theta_1,\theta_2,\tilde{q}(\theta_1,\theta_2)))]],$$

since $\pi(\theta_1, \theta_2, \mu) = 0$ due to the assumption that F is concentrated in $[\theta_1, \theta_2]$. We have thus demonstrated that a price exists which opens the market and is at least as profitable as closing the market.

For the sufficiency of the second claim, without loss in generality, suppose profits are nonnegative at a bid-ask pair where $a \in (\theta_1, \theta_2)$. If the bid side is closed under this pair, resulting in no gains or losses in trade on the bid side, then the ask side will be open under competition at this ask or lower. If the bid side is open and profitable, then the bid side will be open under competition at this bid or greater. If the bid side is open and breakeven or unprofitable under this pair then the ask side must be generating nonnegative profits to satisfy the nonnegativity assertion for the (not necessarily competitive) pair. Thus the ask side of the market will be open under competition at this ask or lower. For necessity, without loss in generality assume $a^c \in (\theta_1, \theta_2)$. Then by definition profits are nonnegative at $b = \theta_1$, $a = a^c$, which is a pair satisfying the required condition.■

Proof of Proposition 5:

The necessity of the condition on F follows from Proposition 4. That such conditions may arise is easily demonstrated by an example. Let $\theta_1 = \frac{1}{3}$ and $\theta_2 = \frac{2}{3}$, F be uniform on [0,1] and consider an arbitrary $n \geq 1$, $\mu \in [0, 1]$, and $\lambda \in (0, 1)$. Since $\pi(\frac{1}{3}, \frac{2}{3}, \mu) > 0$ competitive dealers will open the market at some higher bid and lower ask. Note that this argument does not depend on μ or λ. Therefore the competitive dealers will open the market in all of the T periods. We must show that a monopolist will not open the market in any of the T periods. Consider the final period with arbitrary posterior μ_T. The optimal bid and ask prices are:

$$b_T = \frac{1}{3} - \frac{1}{6}\mu_T \text{ and } a_T = \frac{5}{6} - \frac{1}{6}\mu_T$$

which are less than or equal to $\theta_1 = \frac{1}{3}$ and greater than or equal to $\theta_2 = \frac{2}{3}$, respectively for $\mu_T \in [0, 1]$. Note that λ does not appear since prices are in the range where no informed trade occurs. Therefore, in the T^{th} period, the markets are closed for all possible values of $\mu_T \in [0, 1]$ and $\lambda \in (0, 1)$. Since no information produced prior to period T will be used in period T, the monopolist will myopically optimize in period $T - 1$. But this implies the same price as in period T since the result held for all $\mu_T \in [0,1]$. Thus the market is closed in period $T-1$ and in all previous periods. ∎

References

Admati, A. R., and P. Pfleiderer, 1988, "A Theory of Intraday Trading Patterns," *Review of Financial Studies*, 1, 3–40.

Amihud, Y., and H. Mendelson, 1980, "Dealership Market: Market Making with Inventory," *Journal of Financial Economics*, 8, 31–53.

Easley, D., and N. M. Kiefer, 1988, "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56(5), 1045–1064.

Easley, D., and M. O'Hara, 1987, "Price, Quantity, and Information in Securities Markets," *Journal of Financial Economics*, 19, 69–90.

Gammill, J., 1990, "The Organization of Financial Markets: Competitive Versus Cooperative Market Mechanisms," Working paper, Harvard Business School.

Glosten, L., 1989, "Insider Trading, Liquidity, and the Role of the Monopolist Specialist," *Journal of Business*, 62, 211–235.

Glosten, L., and P. Milgrom, 1985, "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents," *Journal of Financial Economics*, 14, 71–100.

Grossman, S. J., R. E. Kihlstrom, and L. J. Mirman, 1977, "A Bayesian Approach to the Production of Information and Learning by Doing," *Review of Economic Studies*, 44(3), October, 533–547.

Handa, P., 1991, "Order Flow and Bid-Ask Dynamics: An Empirical Investigation," Working paper, New York University.

Harris, L., 1986, "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns," *Journal of Financial Economics*, 16, 99–118.

Harris, L., 1989, "A Day-End Transaction Price Anomaly," *Journal of Financial and Quantitative Analysis*, 24, 29–46.

Hasbrouck, J., 1988, "Trades, Quotes, Inventories and Information," *Journal of Financial Economics*, 22, 229–252.

Hasbrouck, J., 1991, "Measuring the Information Content of Stock Trades," *Journal of Finance*, 46, 178–208.

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

Leach, J. C., and A. Madhavan, 1992, "Intertemporal Price Discovery by Market Makers: Active versus Passive Learning," Forthcoming, *Journal of Financial Intermediation.*

Madhavan, A., 1992, "Trading Mechanisms in Securities Markets," *Journal of Finance*, 47, 500–500.

Madhavan, A., and S. Smidt, 1991, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, 99–134.

Madhavan, A., and S. Smidt, 1992, "An Analysis of Daily Changes in Specialist Inventories and Quotations," Working paper, Wharton School, University of Pennsylvania.

Pagano, M., and A. Röell, 1990, "Trading Systems in European Stock Exchanges," Discussion paper 75, Financial Markets Group, London School of Economics.

Schreiber, P., and R. Schwartz, 1985, "Efficient Price Discovery in a Securities Market: The Objective of a Trading System," in Y. Amihud, T. Ho, and R. Schwartz (eds.), *Market Making and the Changing Structure of the Securities Industry*, Lexington Books, Lexington, MA.

Stoll, H., 1989, "Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests," *Journal of Finance*, 44, 115–134.

Stoll, H., and R. Huang, 1991, "Major World Equity Markets: Current Structure and Prospects for Change," Working paper, Vanderbilt University.

Notes

[1]See e.g., Pagano and Röell (1990) and Stoll and Huang (1991) for a discussion of the issues concerning the merits of alternative market structures; Madhavan (1992) provides a theoretical analysis of different market designs.

[2]See, e.g., Schreiber and Schwartz (1985), who discuss the importance of price discovery in securities markets.

[3]In a related departure from existing models, Leach and Madhavan (1992) derive conditions under which price experimentation will not occur.

[4] While our model addresses short-lived information, the intuition applies to longer-lived information as well. We thank the referee for suggesting this interpretation to us.

[5]We discuss the implications of allowing variable order size below. Leach and Madhavan (1992) provide a model of price discovery where order quantity is variable but traders arrive in batches rather than sequentially.

[6]We do not restrict our analysis to the case where all traders initially have identical expected liquidation values, although our analysis covers this case.

[7]This assumption can be relaxed to include cases where there is a possibility of retrade in future rounds. The current actions are the best replies in this extended model provided the probability of retrade is sufficiently low, especially if the traders' horizon is relatively short. In this case, attempting to manipulate the specialist's beliefs (e.g., by buying when the security is overvalued) is infeasible given time constraints or uneconomical because of competition from other traders or the inability to recoup the losses from misleading the specialist.

[8]The analysis is closely linked to models of a monopolist facing an unknown demand curve (see, e.g., Easley and Kiefer (1988)) and of an agent consuming a good with unknown utility (see, e.g., Grossman, Kihlstrom, and Mirman (1977)).

[9]To simplify the analysis, we assume the discount factor is unity. This assumption has no effect on our qualitative results and is consistent with our motivation of trading on short-lived information.

[10]This distinction is also addressed in Leach and Madhavan (1992) where we consider the relationship of the functional form of *excess* demand to active learning in batch trading systems.

[11]Of course, this could be partially offset by an increased possibility of gaming from the informed if multiple trades by a specific informed trader are possible.

[12]See, e.g., Glosten and Milgrom (1985), Kyle (1985), Easley and O'Hara (1987), and Admati and Pfleiderer (1988).

[13]This assumption is readily modified to allow liquidity traders to update rationally so that the mean of reservation prices is a function of the prevailing bid and ask quotes. This complicates the algebra without altering the qualitative nature of our conclusions.

[14]While our argument relates to a two-state distribution, our intuition about increasing the information content of trade by widening the bid-ask spread follows from the structure of uninformed trade and would also apply to a multi-state distribution.

[15]Of course, some securities markets (re)open trading (like the NYSE in the morning) with a call auction, possibly to facilitate price discovery at lower cost than through specialist price experimentation.

[16]Many studies using intraday data find only weak evidence of inventory control effects. See, e.g., Hasbrouck (1988), Stoll (1989), and Madhavan and Smidt (1991).

[17]Interestingly, Hasbrouck (1991) does find evidence that midpoints do not follow a martingale.

[18]One advantage of using these data is that we do not need to sign volume to determine whether a transaction was at the bid or the ask.

[19]It is possible that an exogenous increase in the asset's volatility could widen competitively determined bid-ask spreads without any price experimentation by the specialist. We attempt to control for this factor by the inclusion of market depth, a proxy for the competition from the limit order book.

[20] The NYSE generally opens with a call auction, following which the system switches to a specialist system. For this reason, we used the first specialist bid-ask quotations in the opening interval to measure the spread.

[20] The NYSE generally opens with a call auction, following which the system switches to a specialist system. For this reason, we used the first specialist bid-ask quotations in the opening interval to measure the spread.

[21]The ordinary-least squares estimates, while significantly stronger than the reported results, exhibited significant positive first-order serial correlation; the Durbin-Watson statistic was only 0.92 on average. Once we had corrected for first-order serial correlation, the Durbin-Watson statistics increased to 2.1 on average, with the smallest being 1.96, and no further adjustments proved necessary.

[22]Competitive pressures are not likely to explain our findings; the opening interval is typically one of the most active periods of the day when the specialist faces the greatest competition.

[23] For example, options on the Chicago Board of Options Exchange (CBOE) are traded in a multiple dealer system, while the American Stock Exchange (AMEX) relies on a specialist system.

Table 1
Variation in Intraday Bid-Ask Quotations

Yule-Walker estimates (with t-values in parentheses) of the regression equation:

$$s_t = \beta_0 + \beta_1|r_{t+1}| + \beta_2|r_t| + \beta_3|I_t| + \beta_4 D_t + \beta_5 O_t + \beta_6 C_t$$

where s_t is the percentage bid-ask spread, I_t is inventory, r_t is the change in the mid-point of the spread from hour $t-1$ to t, D_t is the average depth, and O_t and C_t are dummy variables for the opening and closing intervals, respectively, for the period February–December, 1987, for 16 specialist stocks ranked by transaction frequency from highest to lowest.

Stock	β_0	β_1	β_2	$\beta_3 \times 10^6$	$\beta_4 \times 10^4$	β_5	β_6	R^2
1	1.686	0.250	0.506	4.523	−1.220	0.038	−0.040	0.54
	(23.146)	(2.056)	(4.155)	(5.372)	(−3.237)	(0.947)	(−0.976)	
2	2.567	−0.050	2.787	0.295	−0.195	0.321	−0.039	0.71
	(8.968)	(−0.153)	(8.476)	(2.141)	(−0.741)	(3.624)	(−0.436)	
3	1.785	0.208	0.324	−0.514	−0.396	0.107	−0.054	0.54
	(25.891)	(3.398)	(5.316)	(−7.159)	(−4.051)	(2.946)	(−1.490)	
4	1.814	0.072	0.061	2.653	−0.433	0.200	−0.172	0.48
	(13.740)	(0.565)	(0.792)	(6.650)	(−3.339)	(2.254)	(−1.933)	
5	1.550	0.249	0.702	0.299	−0.152	0.210	−0.078	0.38
	(24.459)	(1.980)	(5.642)	(0.788)	(−2.841)	(4.197)	(−1.559)	
6	0.585	−0.160	1.259	0.767	−0.003	0.162	−0.082	0.47
	(2.292)	(−0.597)	(4.676)	(6.749)	(−0.034)	(2.332)	(−1.176)	
7	0.564	0.040	−0.014	0.433	−0.572	0.041	−0.019	0.64
	(5.683)	(1.528)	(−0.538)	(11.138)	(−4.077)	(0.966)	(−0.446)	
8	1.134	−0.007	−0.010	1.228	−0.492	0.057	0.006	0.61
	(22.210)	(−0.333)	(−0.476)	(13.335)	(−3.924)	(1.584)	(0.176)	
9	0.138	0.205	0.291	0.666	−0.417	0.308	−0.026	0.54
	(1.170)	(2.628)	(3.702)	(16.883)	(−2.979)	(4.963)	(−0.424)	
10	0.562	0.120	0.353	0.061	0.062	0.031	−0.090	0.34
	(21.754)	(3.750)	(11.148)	(2.953)	(6.531)	(1.067)	(−3.067)	
11	1.112	0.509	1.041	1.797	−0.173	0.239	−0.003	0.38
	(19.254)	(3.346)	(6.892)	(11.901)	(−1.725)	(4.427)	(−0.059)	
12	0.615	0.065	0.184	0.117	−0.293	0.043	−0.013	0.52
	(19.465)	(2.889)	(8.231)	(3.734)	(−5.151)	(2.233)	(−0.662)	
13	1.414	0.450	0.552	0.813	−0.071	0.055	−0.019	0.33
	(23.694)	(2.944)	(3.634)	(4.229)	(−0.742)	(1.241)	(−0.416)	
14	0.992	0.090	0.153	0.198	−0.153	0.189	−0.020	0.29
	(15.486)	(1.949)	(3.284)	(4.113)	(−2.456)	(4.901)	(−0.517)	
15	0.816	0.135	0.266	0.198	−0.104	0.074	0.008	0.55
	(29.133)	(3.754)	(7.520)	(22.539)	(−2.749)	(2.137)	(0.217)	
16	0.968	0.426	0.908	0.912	−0.016	0.049	−0.016	0.53
	(18.191)	(4.073)	(8.736)	(14.971)	(−0.612)	(0.971)	(−0.323)	

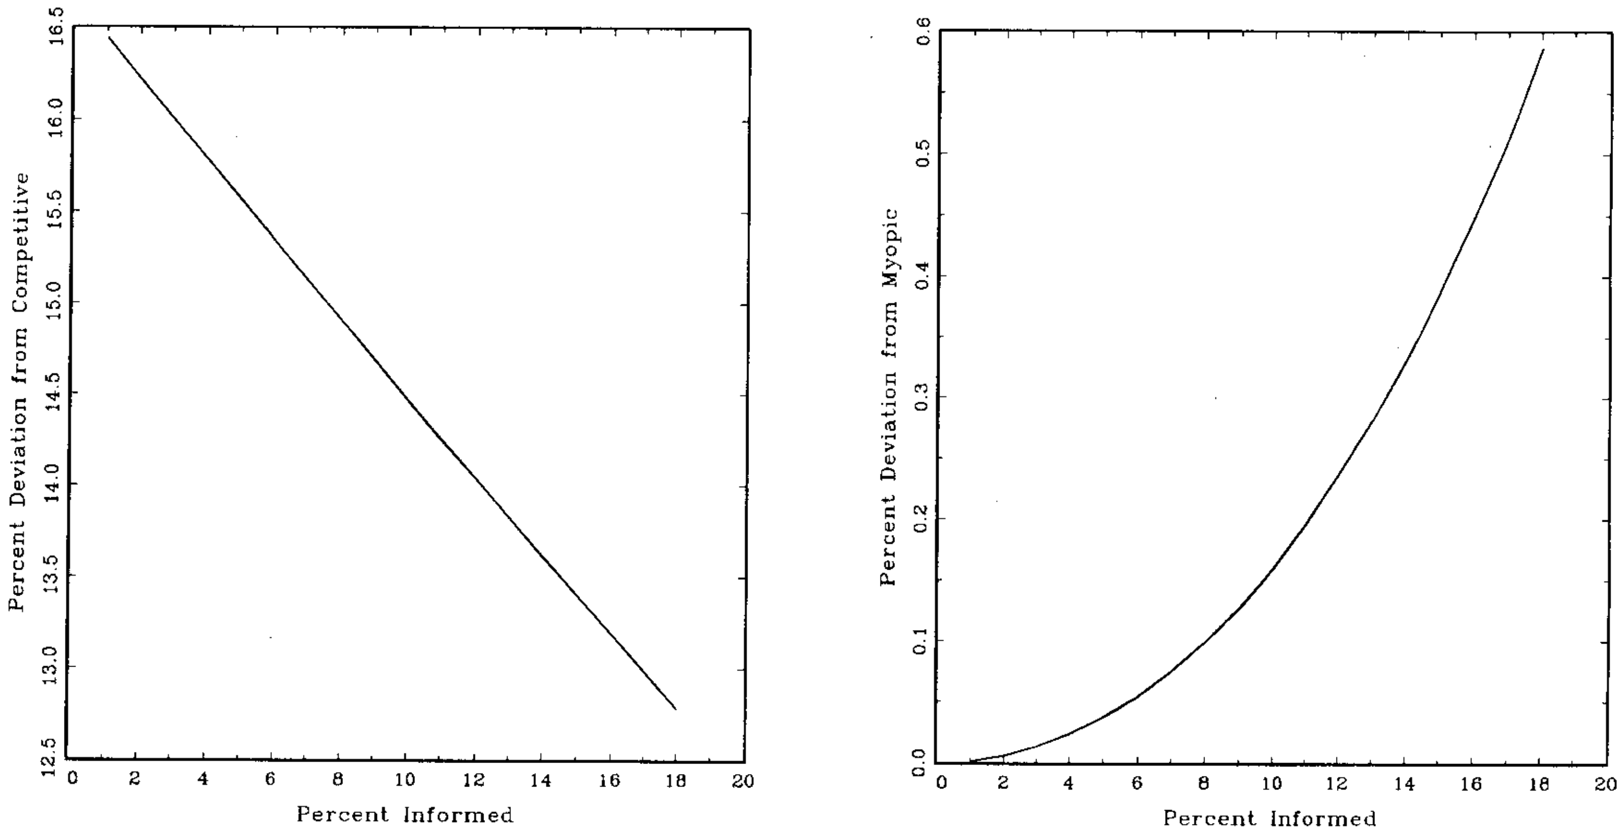

Figure 1
Percentage deviation in ask price from competitve and myopic levels

These graphs display the percentage deviation of a specialist's optimal ask price from the benchmarks of competition and myopia. The graphs are generated using $\Theta_1=1$, $\Theta_2=2$, $\mu=.5$ and F uniformly distributed on the interval [1,2].

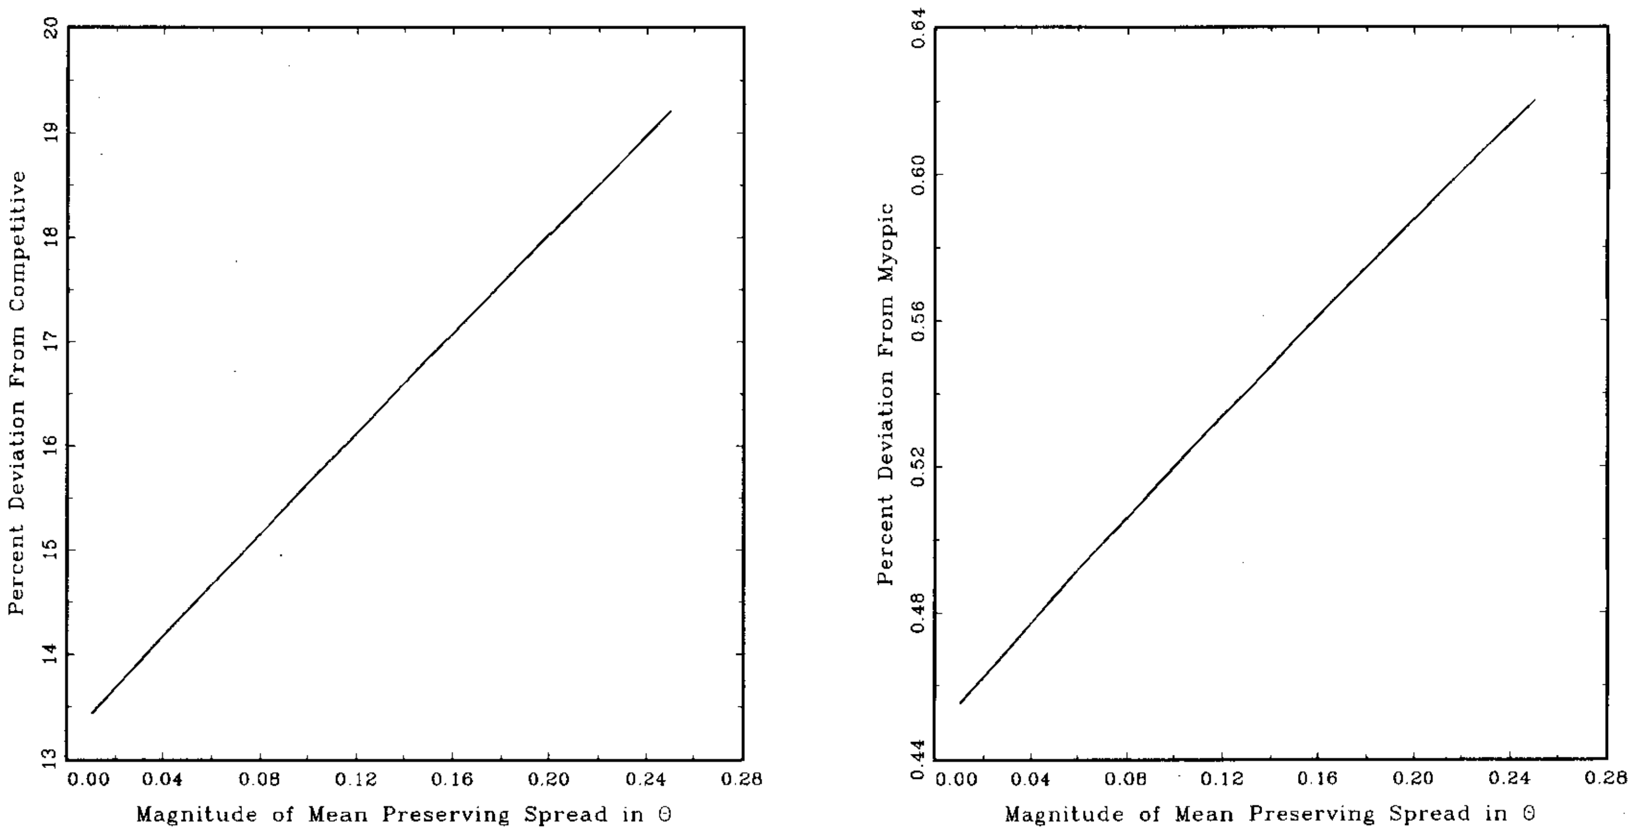

Figure 2
The relationship of experimental pricing to risk

These graphs display the percentage deviation of a specialist's optimal ask price from the benchmarks of competition and myopia as risk in the full−information value increases. The graphs are drawn using $\Theta_1=1$, $\Theta_2=2$, $\mu=.5$ and F uniformly distributed on the interval [1,2].