JOURNAL OF FINANCIAL INTERMEDIATION 2, 207–235 (1992)

Intertemporal Price Discovery by Market Makers: Active versus Passive Learning

J. CHRIS LEACH AND ANANTH N. MADHAVAN*

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Received August 29, 1991

This paper demonstrates that market makers have the ability and incentive to facilitate price discovery in securities markets. Market makers can expedite the process of intertemporal price formation by setting prices to induce statistically more informative order flow. Such actions constitute an investment in the production of information. Under certain conditions, market makers can recoup the cost of this investment by making better pricing decisions in the future with more precise information. These conditions are analyzed in a general model and several examples that illustrate the complex nature of price discovery are presented. *Journal of Economic Literature* Classification Numbers: D82, D83, G12, and G14. © 1992 Academic Press, Inc.

1. INTRODUCTION

Many securities markets rely upon intermediaries, known as market makers or dealers, to provide continuous trading by standing ready to buy or sell securities on demand. These agents have been the subject of numerous studies in finance because of their crucial importance in the provision of liquidity. Early research (e.g., Demsetz, 1968) viewed the market maker as a supplier of "immediacy" and the bid–ask spread as the price charged for this service. To accommodate transitory order imbalances, dealers must take on unwanted inventory. As a result, market makers

* We thank Michael Brennan, David Easley, Richard Kihlstrom, Paul Laux, and two anonymous referees for their valuable suggestions. We have also benefitted from the helpful comments of seminar participants at the American Finance Association Meetings, the Western Finance Association Meetings, Stanford University, and Wharton. Any errors are entirely our own. Financial assistance from the Geewax-Terker Research Fund is gratefully acknowledged.

1042-9573/92 $5.00
Copyright © 1992 by Academic Press, Inc.
All rights of reproduction in any form reserved.

seeking to minimize inventory carrying costs (as in Ho and Stoll (1983)) or to maximize expected utility under risk aversion (as in O'Hara and Oldfield (1986) and Grossman and Miller (1988)) find it optimal to affect not only the size of the bid–ask spread, but also its placement.

More recent models of market making (including Glosten and Milgrom (1985), Kyle (1985), Easley and O'Hara (1987), and Admati and Pfleider (1988), among many others) emphasize the importance of asymmetric information. In these models, the presence of traders with private information implies that order flow conveys a noisy signal regarding the fundamental value of the security. Consequently, rational market makers update their beliefs, and hence prices, in response to order flow. Further, even in the absence of any inventory effects, a bid–ask spread arises because informed traders impose adverse selection costs on market makers. Thus, the asymmetric information models, like the inventory control models, imply that market maker behavior affects both spreads and prices.

The models described above provide important insights into the market maker's dealership function, but they do not shed much light on the role of the market maker in price formation. Market makers simply serve as a mechanism to absorb excess demand at a price that equals the conditional expected value of the security. Thus, the market maker's role in price formation is passive in nature. Yet, because of their central position in trading, it appears natural for market makers to have a crucial role in determining price, especially when there is uncertainty regarding asset value. This paper develops a model in which the market maker actively assists in *price discovery*, i.e., the process by which prices converge to the full-information value of the asset.[1] This role for the market maker emerges as a by-product of intertemporal profit maximization, not as an exchange-mandated function.

We demonstrate that market makers have the ability and the incentives to undertake costly price discovery by experimenting with prices ("testing the waters") to induce statistically more information order flow. This strategy leads to a faster convergence of beliefs and hence more accurate pricing in the future. Intuitively, if different price choices elicit differential order flow responses from informed traders and liquidity-motivated traders, the market maker can set prices to affect the information content (i.e., the signal–noise ratio) of order flow to learn more rapidly about the unobserved market-clearing price. The market maker invests in information production by a costly deviation from the single-period optimal price

[1] Schreiber and Schwartz (1986) note that price discovery is the most crucial function of a market mechanism. Bronfman and Schwartz (1990) and Handa and Schwartz (1989) provide models of price discovery. Unlike our model, these approaches assume no information asymmetry, focusing rather on the effect of the trading mechanism and order form on prices.

choice. The investment cost can be recovered because more precise information enables the market maker to make better decisions in the future. We examine a general model and then construct a series of examples to illustrate our results.

The model differs from market microstructure models such as that of Kyle (1985) in two important respects. First, trading is accomplished here through a single market maker (i.e., a specialist) as opposed to a competitive group of dealers. Second, the market maker sets the price in our model prior to observing order flow, unlike other models in which orders arrive first.[2] The first assumption implies that prices are not constrained to be equal to expected values and that there is an opportunity for the market maker to recoup the expected losses incurred in facilitating price discovery. The second assumption is necessary for the market maker's actions to affect the distribution of order flow.

The main contribution of this paper is to demonstrate that market makers have the ability to play a crucial role in price discovery as noted above. However, we also investigate the optimality of these strategies, showing that price discovery is not optimal for certain functional forms of demand where price does not affect the informativeness of the signal conveyed by order flow. We illustrate the complex nature of price discovery with several examples. We provide an example in which price discovery occurs when there is no asymmetric information concerning fundamental values. In this case, price discovery occurs because it is in the market maker's interest to learn about traders' elasticities of demand in order to maximize his future trading profits. However, in the important case where traders have negative exponential expected utility functions and the future value of the asset is normally distributed (Grossman, 1976), market makers do not have the incentive to engage in price discovery, even though trading reveals information about the private information of traders. We also show that price discovery can affect the bid–ask spread, market maker's profits, and the variability of inventory movements. With a numerical example, we show that price discovery may affect returns and that the amount of price discovery is related in a nonmonotonic fashion to uncertainty regarding fundamentals and to the level of liquidity-motivated trading.

These results have implications for public policies concerning the structure of securities markets because they provide a previously unrecognized benefit to a specialist system with a single market maker. They also have implications for event studies, since the distribution of returns prior to known information events may reflect abnormalities associated with

[2] In this respect, the model resembles quote-driven systems such as the system in operation on the New York Stock Exchange after the opening.

price discovery. We also contribute to a broader literature, since our general analysis can be applied to other problems in securities markets where agents have incentives to learn the underlying parameters of the model, such as in the case of an initial public offering.

The idea of price experimentation as a method to facilitate price discovery is also discussed by Leach and Madhavan (1991). There we examine the differences in price discovery between a specialist system and a multiple dealer system in a model where order size is fixed and market makers experiment by changing the size of the spread. Similarly, Gammill (1990) examines market structures in which dealers knowingly trade with an informed trader in order to "buy" her private information. In the model presented here, however, the market maker cannot identify with certainty an informed trader. Further, our focus in this paper is not on a comparison of price formation across different market structures, but on the conditions necessary for price discovery to occur in the first instance. Price discovery is a complex phenomenon and a formal description of the market maker's optimal intertemporal pricing behavior is needed to illustrate the subtleties that lie behind what appears intuitively to be a simple strategy.

The model presented here is most closely related to the experimental consumption model of Grossman *et al.* (1977) (hereafter GKM), in which agents, who are uncertain about their utility from consuming a good, experiment to obtain information to improve their consumption decisions in the future.[3]

Although the problem we analyze is very similar to the experimental consumption model, we cannot apply the results from this literature directly to securities markets because the experimental consumption literature typically assumes that the agent observes her utility immediately upon consumption of the good. In contrast, the market maker's reward cannot be observed after the current period's action because utility depends on changes in expected terminal wealth due to current actions, and terminal price is unobservable. Further, we are concerned with characterizing the conditions under which price discovery does or does not occur, so we state our results for more general functional forms.

The remainder of the paper proceeds as follows. Section 2 develops our model and examines the conditions under which market makers have the incentive and the opportunity to facilitate price discovery. Section 3 provides examples to illustrate the complex nature of price discovery. The

[3] See also Kihlstrom *et al.* (1984). The present analysis also resembles models of a monopolist facing an unknown demand curve (e.g., McLennan, 1984; Easley and Kiefer, 1988) and to bandit problems (see Rothschild, 1974; Berry and Fristedt, 1985). We differ from these models in that our focus is on the optimality of experimental behavior rather than the asymptotic properties of beliefs.

examples show that price discovery can affect asset prices and that there can be a marked departure from the predictions of single-period models. In Section 4, we discuss the implications of our results for empirical research and for public policy. Finally, Section 5 concludes the paper. All proofs are contained in the Appendix.

2. THE MODEL

We consider the intertemporal pricing decision following an information event which induces a shift in the demand or supply of a risky asset. We assume that occurrence of the information event is common knowledge and that the monopolistic market maker knows the functional form, but not the parameters, of the excess demand curve. Trading occurs over a finite number of trading rounds, indexed by $t = 1, \ldots, T$. The final date T is a deterministically known information event when uncertainty is resolved and the security pays a liquidating dividend. Thus, the model is best viewed as a description of trading preceding an event such as a dividend or earnings announcement the timing of which is known.

The trading protocol is as follows: In trading round t the market maker quotes a single price, denoted p_t, and then observes the excess demand at that price, denoted $z(p_t)$. After observing an excess demand realization, he updates his beliefs about the parameters of excess demand and sets a price for the next round. The assumption of a single price greatly simplifies the exposition and can be relaxed as we show in Section 3.3. We begin by considering a general specification for demand; in Section 3.1 below we provide an example where we explicitly describe the optimization problem that gives rise to such a function. Excess demand in round t takes the form

$$z_t = f(p_t, \beta) + \varepsilon_t, \tag{1}$$

where $f(p_t, \beta)$ is the price-sensitive portion of excess demand which we assume to be differentiable and strictly decreasing in p_t. The function f also depends upon β, the unknown parameter, which is drawn randomly from $B = \{\beta^1, \ldots, \beta^n\} \subset \Re^n$.[4] Finally, ε_t is the stochastic disturbance term which we assume to be normally distributed with $E[\varepsilon_t] = 0$; $E[\varepsilon_t \varepsilon_{t'}] = 0$ if $t \neq t'$ and σ^2 if $t = t'$.[5] Let $h(\cdot; \mu)$ denote the normal density function

[4] Here β could be regarded as a vector of coefficients, where one element of the vector is unknown. We chose to consider a finite sample space for β for convenience when constructing examples. Our arguments do not require a discrete set of β's.

[5] In what follows the additivity of the error term is inessential and we chose this as an analytic convenience. Our derivations do, however, rely on lack of correlation in β and ε_t and independence in ε_t's.

with mean μ and variance σ^2. When $F(\cdot) \equiv f(\cdot; \beta)$ is invertible, we can define $\theta \equiv F^{-1}(0)$. The constant θ is interpreted as the full-information price, i.e., the price that is expected to clear the market in the absence of any extrinsic uncertainty.

The market maker knows this structure, but he does not know the realization of β, and hence the market-clearing price, θ. Each period, the market maker receives a noisy information signal about β in the form of the observed excess demand at the price set. The market maker's prior probability that $\beta = \beta^i$ in period t is denoted by μ_t^i, where $0 \leq \mu_t^i \leq 1$, for $t = 0, \ldots, T$. It is convenient to define $\mu_t = \{\mu_t^i\}_{i=1}^n$. Boldface variables indicate the time series vectors.

For technical reasons, we need to bound the feasible price space. We assume that the exchange requires that $p_t \in [0, \bar{p}]$, where $0 < \bar{p} < \infty$. Define $I \equiv [0, \bar{p}]$, the domain of p_t. We refer to the set I as the action space. Let $u(p, \beta^i, z)$ be the utility from a trade volume of z at price p when $\beta = \beta^i$. We assume that u is continuous in p.[6] In what follows, we show that the market maker has an incentive to discover full-information value by setting prices in early periods to improve the quality of information revealed by order flow. Such experimentation dictates a departure from the short-run optimum in order to obtain information which is used to increase the expected utility of future consumption. We turn now to a formal statement of the problem.

2.1. *Myopic Optimization and Passive Learning*

The single-period expected utility of the market maker at time t for a price choice of p_t and beliefs μ_t is

$$U(p_t, \mu_t) = \sum_i \mu_t^i \int_{\Re} u(p_t, \beta^i, f(p_t, \beta^i) + \varepsilon_t) h(\varepsilon_t; 0)\, d\varepsilon_t, \tag{2}$$

which can be written (by change of variable) as

$$U(p_t, \mu_t) = \sum_i \mu_t^i \int_{\Re} u(p_t, \beta^i, z_t) h(z_t; f(p_t, \beta^i))\, dz_t. \tag{3}$$

We assume that the structure on f and u is sufficient to generate a unique

[6] For example, the market maker could maximize ex post profits or the utility of trading income. The reward function may also implicitly incorporate a brokerage fee which the market maker charges for the net order imbalance he accommodates. Without such compensation, the market maker would suffer expected losses, because traders would on average sell when prices were set too high and buy when prices were too low. Note, however, this form of the utility function does not allow utility to depend on past realizations or actions other than through beliefs about β.

myopically optimal price.[7] For period t, under prior μ_t, we denote this price

$$p_t^M \equiv \text{argmax}_{p_t} \{U(p_t, \mu_t)\}. \tag{4}$$

We refer to a $p_t \neq p_t^M$ as *active learning* or *price experimentation*. For a rational market maker, departure from p_t^M is only attributable to investment in the production of information. Note that the market maker gathers information about the parameters of the excess demand function, not the security's fundamental value per se. To the extent that the parameters of the excess demand function reflect the unknown fundamental value (e.g., because of the presence of traders with private information), this is not a serious limitation of the analysis.

2.2. *Intertemporal Optimization*

To understand the intertemporal decision process in the dynamic learning environment more clearly, we need to specify the informational link between trading rounds. Following the approach of GKM we let $(\mathbf{z}_t, \mathbf{p}_t) = ((z_0, \ldots, z_t), (p_0, \ldots, p_t))$ for $t = 0 \ldots T$. Then, $(\mathbf{z}_t, \mathbf{p}_t)$ represents the history of trading through period t. Let H_t designate the set of all possible time t histories $(\mathbf{z}_t, \mathbf{p}_t)$. Using the theory of sequential Bayesian updating, we can write Bayes' rule for this problem as a t-period updating problem with the initial prior $\mu_t = g_t(\mu_0, \mathbf{z}_{t-1}, \mathbf{p}_{t-1})$ or as a one-period updating problem with the $t - 1$ prior $\mu_t = \mu_t(\mu_{t-1}, z_{t-1}, p_{t-1})$.

Let $M_t \equiv \{\mu_t : \mu_t = g_t(\mu_0, \mathbf{z}_{t-1}, \mathbf{p}_{t-1}), (\mathbf{z}_{t-1}, \mathbf{p}_{t-1}) \in H_{t-1}\}$, i.e., the set of all possible posteriors that could occur at time t. Finally, let $\delta \leq 1$ be the market maker's discount factor. Now define

$$V_T(\mu_0) \equiv \max_{p_0, \{p_t\}_{t=1}^T} \left\{ E\left[U(p_0, \mu_0) + \sum_{t=1}^{T} \delta^t U(p_t, g_t(\mu_0, \mathbf{z}_{t-1}, \mathbf{p}_{t-1}))\right]\right\}, \tag{5}$$

where $p_0 \in I$, and p_t: $M_t \mapsto I$. In (5), $V_T(\mu_0)$ denotes the discounted sum of expected utility with T periods remaining starting from a prior of μ_0 when Bayes' rule is used to update prior beliefs. By Bellman's optimality principle for p_t: $M_t \mapsto I$ we have

$$V_T(\mu_0) = \max_{p_0} \{U(p_0, \mu_0) + \delta E[V_{T-1}(\mu_1(\mu_0, z_0, p_0))]\}, \tag{6}$$

[7] For example, a sufficient condition for a unique myopic price is that the utility function u is a concave function of profits and the demand function f is a continuous and differentiable function of price p.

where

$$E\{V_{T-1}[\mu_1(\mu_0, z_0, p_0)]\} = \int_{-\infty}^{+\infty} V_{T-1}[\mu_1(\mu_0, z_0, p_0)] \sum_j \mu_0^j h(z_0; f(\beta^j, p_0)]\, dz_0. \tag{7}$$

An optimal price in period 0, given a prior of μ_0, is[8]

$$p_0^* \in P_0^* \equiv \operatorname{argmax}_{p_0} V_T(\mu_0). \tag{8}$$

Equations (6) and (7) are a formal representation of the decision problem facing the market maker. Our model differs from other models of intertemporal price formation because the market maker attempts to change the statistical properties of order flow to induce the revelation of information. We proceed by formalizing this idea.

2.2.1. Characterizing the Information Content of Trade. In this section, we characterize the intuitive notion of more information order flow. This characterization allows a formal analysis of the market maker's price setting behavior using results from statistical decision theory developed by Blackwell (1953). To accomplish these objectives, we need to (1) restate the market maker's objective function in such a way that the maximization problem is seen to be the expectation over β (using the prior distribution) of some function of the actions of the market maker and the unknown parameter β (and not a function of the prior probability distribution); and (2) specify how one price can be more informative than another. We begin with a useful lemma that allows us to restate the market maker's problem in a form in which Blackwell's results can be directly applied:

LEMMA 1. *There exist strategies,* $\xi = (p_0, \gamma)$, *where* $p_0 \in I$, *and* γ_t: $M_{t-1} \mapsto I$ *and functions* $\eta_T(\xi, \beta^i)$ *such that* $V_T(\mu_0)$ *can be written in the form*

$$V_T(\mu_0) = \sup_{\xi} \sum_i \eta_T(\xi, \beta^i)\mu_0^i. \tag{9}$$

The lemma shows that the prior probabilities can be separated from the remainder of the objective function, which allows us to apply Blackwell's main result on comparison of experiments. We can now consider what it

[8] Technically, to ensure V_T is well-defined and P_0^* is nonempty, we need to show that the value function is a continuous function of p, for $p \in I$. This involves showing that $u(\cdot)$ and $f(\cdot; \beta)$ are continuous in p, and $V(\cdot)$ is convex and therefore continuous in μ (on open convex sets). For a prototype of the proof technique, see, e.g., GKM or Prescott (1972).

means in our model to say that one price is more informative than another. Recall that $h(z; f(p, \beta))$ represents a normal density function of z with mean $f(p, \beta)$ and variance σ^2, which we write as $h(z; p, \beta)$. As the action variable p varies, we obtain a family of distributions or experiments, $\{h(\cdot; p, \cdot), p \in I\}$. Distinct elements of this set correspond to different experiments. We refer to the experiment p by which we mean $h(\cdot; p, \cdot)$. Our definition of a more informative experiment is the sufficiency condition of Blackwell (1953), which is also discussed by DeGroot (1970).

DEFINITION 1. An experiment p is more informative than (or is sufficient for) an experiment p' if there exists a real-valued function $\nu(z'; z) \geq 0$ such that for all $z' \in \Re$ and all $\beta \in B$,

$$h(z'; p', \beta) = \int \nu(z'; z)h(z; p, \beta)\, dz, \tag{10}$$

where $\int \nu(z'; z)\, dz = 1$.

The basic idea of sufficiency is that the results of the experiment p' can be replicated (in a probabilistic sense) by auxiliary randomization (density ν) given the results of the experiment p. Therefore, p is more informative than p'.

By restating the market maker's optimization problem and introducing the concept of more informative prices, we can apply Blackwell's lemma on the comparison of experiments to characterize the market maker's pricing strategy:

LEMMA 2 (Blackwell). *Consider functions of the form*

$$V_*(\mu) = \sup_{\xi \in \nabla} \sum_i \eta_T(\xi, \beta^i)\mu^i, \tag{11}$$

where the set ∇ represents the set of possible decisions, an element of which is designated ξ, and η: $\nabla \times B \rightarrow \Re$. If the experiment with p is more informative than the experiment with p' then the expected value of a function in this V_* class is higher under p than under p',

$$E[V_*(g_1(\mu, p, z))] \geq E[V_*(g_1(\mu, p', z'))], \tag{12}$$

where

$$E[V_*(g_1(\mu, p, z))] \equiv \int_{\Re} V_*(g_1(\mu, p, z)) \sum_i h(z; \beta^i, p)\mu^i\, dz \tag{13}$$

and $E[V_*(g_1(\mu, p', z'))]$ is similarly defined.

Blackwell's lemma states that more informative prices give rise to higher expected future values. By Lemma 1, we know that $V_{T-1}(\mu_1)$ is in the class of $V_*(\mu)$ functions where

$$E[V_{T-1}(m_1(\mu_0, z_0, p_0))] = \int_{\Re} V_{T-1}(m_1(\mu_0, z_0, p_0)) \sum_i h(z; \beta^i, p)\mu_0^i \, dz. \tag{14}$$

Using Lemma 2, we obtain

$$E[V_{T-1}(g_1(\mu_0, z_0, p_0))] \geq E[V_{T-1}(g_1(\mu_0, z_0', p_0'))] \tag{15}$$

when the p is more informative than p'. Thus, more informative experiments give rise to higher future expected profits. The intuition for our result is straightforward; actions that raise the signal-to-noise ratio by deviating from the optimal single-period action result in lower current utility to the market maker but allow for more accurate decisions to be made in the future, increasing the expected future profits. This is true for all functional forms for excess demand and is independent of the discount factor, δ.

In summary, these results in statistical decision theory make precise the intuitive idea that a price which induces a more informative order flow increases future expected utility. While this characterization tells us that more informative current prices give higher future utility, it does not indicate when future gains will outweigh current (opportunity) losses incurred in departing from myopically optimal pricing. Therefore, we digress to characterize situations in which myopic prices are still optimal, i.e., $p_t^M = P_t^*$.

2.2.2. Optimal Myopic Pricing—Additively Separable f. Suppose that the realization of β is β^i and $f(p_t, \beta^i)$ is such that the path of posterior beliefs regarding β is invariant to the price choice. In this case, the market maker has no ability to control the learning process, i.e., to affect the information content of the signals conveyed by order flow. In our problem, because order flow shocks are uncorrelated with β, all future beliefs will be independent of the current price choice if and only if the next period's beliefs are independent of the current price choice. Formally,

LEMMA 3. *If $p_t = p(\mu_t, t)$ for all t, then $\boldsymbol{\mu}_s(\mathbf{z}_{s-1}(\mathbf{p}_{s-1}), \mathbf{p}_{s-1}, \mu_0)$ is independent of p_0 for all $s > 0$ if and only if $\mu_1(z_0(p_0), p_0, \mu_0)$ is independent of p_0.*

A market maker who does not know the true β will know that attempting to control the learning process is futile if and only if next period's beliefs do not vary with current price for any $\beta \in B$. When this invariance holds, it is unnecessary to look beyond one period into the future, as long as actions depend only upon the state variables for this problem, i.e., current beliefs and the current time period. Proposition 1 establishes conditions under which the market maker has no incentive to deviate from the single-period optimal price choice.[9]

PROPOSITION 1. *Learning is independent of current price if and only if the expected order flow, $f(p, \beta)$, satisfies*

$$\frac{\partial f(p, \beta)}{\partial p} = \psi(p)$$

for all $p \in I$, where $\psi(p)$ is a real-valued function which is independent of β. If learning is independent of current price, the optimal dynamic actions coincide with the myopic action (i.e., $P_t^* = p_t^M$).

The proposition implies that if there is no interaction between price and the parameter β, the market maker knows that his actions are not going to alter the statistical information received in the next period. Accordingly, if f has this property, then the optimal price choice in any period is a single-period or myopic utility maximizing choice given current information. Departing from myopic optimization does not refine the information from trade, and an optimizing market maker will not incur any current costs associated with such a departure. For functions failing to satisfy this criterion, controlled learning opportunities exist for at least one value of β. All β's occur with positive probability and, therefore, the market maker knows that prices can alter future beliefs. However, this alone does not imply that the market maker departs from myopic pricing, only that such a departure might prove beneficial.

2.2.3. Departure from Myopia: Multiplicatively Separable f. Proposition 1 shows that pricing strategies for functions $f(p, \beta)$, where p and β are additively separable, do not permit controlled learning. For functions that do not have this property, opportunities arise for active learning through price experimentation. When will such active learning strategies be optimal and what predictions can we make about price behavior under such strategies? The answer to this question will typically depend on the functional form of f, which in turn depends upon the behavior of traders in the market. Proposition 1 eliminates additively separable functions as candidates for active learning, so we turn to multiplicatively separable functions.

[9] We state our results using our assumption that f is differentiable in p. The proof is easily generalized to apply to functions which are not differentiable.

Let $f(p, \beta) = \psi_0 + \psi_1(\beta)\psi_2(p)$, so that excess demand is given by

$$z_t = \psi_0 + \psi_1(\beta)\psi_2(p) + \varepsilon_t. \tag{16}$$

We assume, without loss of generality, that $\psi_1 \neq 0$ takes on at least two values and that $\psi_2 > 0$.[10] The assumptions that f is differentiable and downward sloping imply that ψ_2 is differentiable and $\psi_1(\beta)\psi_2'(p) < 0$. The nondegeneracy of ψ_1 and the strict negativity of $\psi_1(\beta)\psi_2'(p)$ imply that Proposition 1 is inapplicable because $\partial f/\partial p = \psi_1(\beta)\psi_2'(p) < 0$ depends on β. We now establish a useful result for demand functions of the type considered above.

LEMMA 4. *For multiplicatively separable excess demand functions of the form* (16), *the order flow induced by p is more informative than that induced by p' if $\psi_2(p) \geq \psi_2(p')$.*

The intuition for the result may be demonstrated using signal-to-noise ratios. Consider the experiment with p. Observing z is equivalent to observing $r = (z - \psi_0)/\psi_2(p) = \psi_1(\tilde{\beta}) + \tilde{\varepsilon}/\psi_2(p)$, which is normally distributed with unknown mean $\psi_1(\tilde{\beta})$ and variance $\mathrm{Var}(\psi_1(\tilde{\beta})) + \sigma^2/\psi_2(p)^2$. The (modified) signal-to-noise ratio is

$$\frac{\psi_1(\tilde{\beta})}{\mathrm{Var}(\psi_1(\tilde{\beta})) + \sigma^2/\psi_2(p)}$$

which is decreasing in $\psi_2(p)$. Higher ψ_2's, while not affecting the signal content of order flow, do decrease the noise content. So, the higher the $\psi_2(p)$, the more informative the experiment. The following proposition formally establishes this result:

PROPOSITION 2. *Optimal prices induce higher values of ψ_2 than myopic prices. Specifically, define $\psi_2^* = \min\{\psi_2(p) : p \in P_0^*\}$ and $\psi_2^M = \psi_2(p_0^m)$. Then $\psi_2^* \geq \psi_2^M$.*

In summary, for multiplicatively additive functional forms of excess demand, an optimizing market maker will always select prices which induce higher levels of ψ_2. However, it need not be the case that the optimizing monopolist raises prices relative to the myopically optimal level. For example, for $\psi_2 = p$, an optimizing monopolist chooses prices above the myopically optimal level, because higher price levels induce more revealing order flow. By contrast, when $\psi_2 = 1/p$, he chooses prices below the myopically optimal level because a lower price level implies a higher signal-to-noise ratio. In general, for the class of multiplicatively

[10] This restriction is just a sign convention that allows us to state our results for the choice variable p in a compact form.

separable f's, the relationship of experimentation to transaction prices and the level of excess demand will depend on the specification of ψ_2's. We can, however, unambiguously compare variability in the market maker's trading under experimentation:

PROPOSITION 3. *For z satisfying Eq.* (16), *the variability in the market maker's change in inventories in the first round of trade* (*i.e.*, z_0) *is greater under p^* than under p^M. Further, if the market maker is risk neutral, the expected profits in the first trading round are lower under p^* than under p^M.*

Intuitively, a deviation from the single-period strategy implies greater variability in order flow and hence in specialist inventories. If the market maker is risk neutral, costly price discovery implies lower expected profits in the first period.

Proposition 3 has some important implications for observed returns around information events. Suppose that an econometrician uses a model of the type $p_t^M(g_t(\mu_0, \mathbf{z}_{t-1}, \mathbf{p}_{t-1})$ to predict prices in round t, where $g_t(\cdot, \cdot)$ is the t-period Bayes operator and $(\mathbf{z}_{t-1}, \mathbf{p}_{t-1})$ is the history of trading. Such a model will work well if information is homogeneous and the optimal and myopic policies coincide. Following an information event, however, the market maker's search for a new settling price distorts prices (and returns) relative to the predictive model. An econometric study that attempts to infer positive or negative stock price reactions using a predictive model based on historical data will also capture price discovery effects. In addition, if the market maker is risk neutral, the gross returns of traders in the early rounds will not reflect the true change in value from the preinformation to the postinformation period.

3. APPLICATIONS

In this section, we develop a series of examples to illustrate the propositions of the previous sections. Our first example shows that active learning can occur even if order flow contains no information about the liquidation value of the security. Price discovery takes the form of learning the market settling price, which depends on investors' time preferences. In the second example, the market maker trades with privately informed agents, but the demand function is such that the market maker has no opportunity to control the learning process. We then extend our results to the case of a two-sided market with an endogenously determined bid–ask spread. Finally, we provide a numerical example to illustrate that active control strategies can strictly dominate myopic strategies. This last example suggests that price discovery can have a nontrivial effect on asset returns.

3.1. *Example 1: Learning Unrelated to Liquidation Value*

Consider a multiperiod model where $N \times T$ investors trade claims to a risky security. Calendar time is divided into 2 "days." On the first day, Day 0, there are T trading rounds. Following trading, the security pays a liquidating dividend at Day 1 in the future. The dividend is distributed normally with mean v_0 and variance σ_v^2. We assume that on each of the T preannouncement trading rounds, N traders arrive at the market through some exogenous arrival process. We assume investors trade only once. An investor $i = 1, \ldots, N$, who trades at round $t = 1, \ldots, T$, has an endowment of cash, $W_{it,0}$, and risky securities, x_{it}. The endowment vector $\omega_{it} = (W_{it,0}, x_{it})$ is private information. Across all investors, x_{it} is drawn from a normal distribution with mean zero and constant variance. Let $C_{i,1}$ represent the cash flow to investor i in the final period, i.e., Day 1. We assume that a trader's subjective future value of income, denoted by C_i, is given by

$$\tilde{C}_i = \rho_i C_{i,0} + \tilde{C}_{i,1}. \tag{17}$$

Investors maximize the expected utility of $\tilde{C}_i$, where for simplicity we assume a mean-variance utility function,

$$u_i(\tilde{C}_i) = E[\tilde{C}_i] - \left(\frac{r}{2}\right) \mathrm{Var}[\tilde{C}_i], \tag{18}$$

where $r > 0$ is the coefficient of absolute risk aversion.[11] The parameter ρ_i is investor i's rate of time preference, where $\rho_i \geq 1$. Future income is a random variable, because it depends on the risky asset's total payoff. For investor i, who trades in round t, let $q_{it} \in \Re$ denote the trade size, with the convention that q_{it} is positive for purchases and negative for sales. The investor's budget constraint implies that we have

$$C_{i,0} = W_{it,0} - p_t q_{it} \tag{19}$$

$$\tilde{C}_{i,1} = \tilde{v}(q_{it} + x_{it}). \tag{20}$$

Then, solving the utility maximization problem, we obtain

$$q_{it} = \frac{E_i[\tilde{v}] - \rho_i p_t}{r\sigma_v^2} - x_{it}. \tag{21}$$

[11] Unlike the usual constant absolute risk aversion utility function, this is a two-parameter utility function.

In what follows, we assume that all investors have the same preference parameter $\rho_i = \rho$. Then, in round t, we can write

$$z_t = \sum_{i=1}^{N} q_{it} = N\left(\frac{v_0 - \rho p_t}{r\sigma_v^2}\right) + \varepsilon_t, \tag{22}$$

where $\varepsilon_t \equiv -\Sigma\, x_{it}$ is distributed normally with mean zero and constant variance. We assume that the market maker does not know the value of ρ, but knows that $\rho \in \{\rho_1, \ldots, \rho_n\}$. Putting $\psi_0 = N_{v_0}/r\sigma_v^2$, $\beta = N_\rho/r\sigma_v^2$, $\psi_1(\beta) = -\beta$, and $\psi_2(p_t) = p_t$, this can be written as

$$z_t = \psi_0 + \psi_1(\beta)\psi_2(p_t) + \varepsilon_t, \tag{23}$$

where $\tilde{\beta} \in \{\beta^1, \ldots, \beta^n\}$ is a random variable from the market maker's perspective. Note that in this example the market-clearing price is distinct from the expected liquidation value; the market-clearing price is $\theta = \psi_0/\beta = v_0/\rho$ while the fundamental value of the asset is $\tilde{v}$. From Eq. (3), the myopic utility for a risk neutral market maker is given by

$$U(p_t, \mu_t) = \sum_i \mu_t^j \int_{\Re}\int_{\Re} (p_t - \tilde{v}) z_t h(z_t; f(p_t, \beta^j)) h_v(v)\, dz_t\, dv, \tag{24}$$

where h_v is the density of the unknown future liquidation value. This can be written as

$$U(p_t, \mu_t) = (p_t - v_0)(\psi_0 - p_t E[\tilde{\beta}]), \tag{25}$$

where v_0 is the expectation of $\tilde{v}$. From the definition of ψ_0 and β, we see that the myopic action is

$$p_t^M = \frac{v_0}{2}\left(1 + \frac{1}{E[\tilde{\rho}]}\right), \tag{26}$$

which is an average of v_0 and the expected clearing price, $E[\theta] = E[v_0/\tilde{\rho}]$. Now consider the active learning price. From Proposition 2, we note that $p_0^* > p_0^M$, i.e., that active learning leads to "overpricing" relative to the passive or myopic learning.[12] This example demonstrates that the myopic and optimal prices diverge even though, in the absence of asymmetric information, trading does not convey a signal regarding the value of the security, $\tilde{v}$. This is because the optimal single-period price depends upon

[12] It is not clear whether there exists a closed-form solution for p_0^*; given parameter values, the solution can be computed numerically as shown in Example 4.

how investors trade off current and future consumption. A market maker who knows the rate of intertemporal price substitution is better able to extract rents from traders. This effect leads to experimentation in the case in which ρ is unknown.

3.2. *Example 2: No Price Experimentation with Asymmetric Information*

Consider a modification of the model above in which each investor receives a private information signal $y_i \in Y$, where Y is a finite set. Suppose y_i and $\tilde{v}$ have known nonzero correlation and $\rho_i = 1$. The market maker does not obtain a private information signal and his information is limited to the publicly available prior signal. In this case, the excess demand is round t is

$$z_t = \sum_i^N \left(\frac{E[\tilde{v}|y_i] - p_t}{r\sigma^2_{\tilde{v}|y}} \right) + \varepsilon_t, \tag{27}$$

where $\varepsilon_t = -\Sigma\, x_{i,t}$, as before, and $\sigma^2_{\tilde{v}|y}$ is interpreted as the conditional variance of $\tilde{v}$ given prior beliefs and the signal. Uncertainty in this model concerns the future value of the asset, $\tilde{v}$. From (27), we can write f as

$$f(p;\, \beta) = \phi\beta - \phi p_t, \tag{28}$$

where we define $\phi \equiv N/r\sigma^2_{\tilde{v}|y}$ and $\beta \equiv \Sigma_i\, E_i[\tilde{v}|y_i]/N$ is the mean conditional expectation of the security's value.[13] The function $f(p_t;\, \beta)$ satisfies

$$\frac{\partial f}{\partial p_t} = -\phi, \tag{29}$$

and applying Proposition 1, no price experimentation will occur because the price set does not affect the posterior distribution of $\tilde{v}$. The optimal dynamic action at any round t is the myopic action. If the market maker is risk neutral, $p_t = E_t[\tilde{v}]$, where the expectation is taken on public information at time t.

3.3. *Example 3: Learning and the Bid–Ask Spread*

The introduction of a two-sided market complicates our model, but is desirable if we wish to discuss the effects of active learning on the bid–ask spread. Suppose the market maker now sets two prices in trading round t,

[13] Since the set Y is finite, β lies in a finite set given by the combinations of all possible conditional expectations.

a bid price, denoted b_t, and an ask price, a_t. All buy orders are executed at the ask price and all sell orders are executed at the bid price. We assume that the supply of securities to the market maker, q_t^s, is given by

$$q_t^s = q^s(b_t; \beta_s) + \varepsilon_t^s, \tag{30}$$

where $q^s(\cdot; \beta_s)$ is increasing in b_t and $\beta_s \in \{\beta_s^1, \ldots, \beta_s^n\}$ is a parameter that is unknown to the market maker. The error term ε_t^s is assumed to be normally independently and identically distributed for all periods.[14] Similarly, the quantity of securities demanded from the market maker, q_t^d, is assumed to be a decreasing function of the ask price, a_t, and we write

$$q_t^d = q^d(a_t; \beta_d) + \varepsilon_t^d, \tag{31}$$

where, again, the error term ε_t^d is assumed to be distributed normally and independently of ε_t^s. The $q^d(\cdot; \beta_d)$ function is assumed to be differentiable and strictly decreasing in a_t. We assume that there exists a unique $\theta > 0$ such that

$$q^d(\theta; \beta_d) = q^s(\theta; \beta_s), \tag{32}$$

i.e., that at a price θ the price-dependent component of demand equals the price-dependent component of supply. The expected profit function in period t is

$$\begin{aligned}\pi(a_t, b_t) &= \sum_i \left\{\int_{\Re}\int_{\Re} [(a_t - \theta)q_t^d \right.\\ &\qquad \left. + (\theta - b_t)q_t^s]\, h(\varepsilon_t^s)\, h(\varepsilon_t^d)\, d\varepsilon_t^s\, d\varepsilon_t^d\right\} \mu_t^i \qquad (33)\\ &= \sum_i \left\{\int_{\Re} [(a_t - \theta)q_t^d]\, h(\varepsilon_t^d)\, d\varepsilon_t^d\right\} \mu_t^i\\ &\qquad + \sum_i \left\{\int_{\Re} [(\theta - b_t)q_t^s]\, h(\varepsilon_t^s)\, d\varepsilon_t^s\right\} \mu_t^i, \qquad (34)\end{aligned}$$

where $\mu_t^i = \Pr[\theta = \theta_i]$ given information at time t. From this we can see that the problem is separable into demand-side and supply-side problems, each of which has an analogous structure to our excess-demand formulation. Therefore our lemmas and propositions can be applied to each side separately. What is the impact of the price discovery on the bid–ask spread? As in the excess demand formulation, the answer depends upon

[14] Unfortunately, this implies that supply could be negative with positive probability. We assume the variance of ε_t^s is sufficiently small relative to supply that the probability of this event is arbitrarily small. Similar comments apply to the demand function.

the specific functional forms taken by the demand and supply curves. Suppose that the demand curve is

$$q^d(a_t; \beta_d) = \psi_0^d + \psi_1^d(\beta_d)\psi_2^d(a_t) \tag{35}$$

and the supply curve is

$$q^s(a_t; \beta_s) = \psi_0^s + \psi_1^s(\beta_s)\psi_2^s(b_t), \tag{36}$$

where $\psi_2^j(\cdot) > 0, j \in \{s, d\}$, is differentiable and nonzero. Proposition 1 (with q identified) with f) implies that active learning opportunities exist when the ψ_1 functions are not single valued. Further, Proposition 2 implies that b is more informative than b' if

$$\psi_2^s(b) \geq \psi_2^s(b'). \tag{37}$$

Similarly, the experiment a is more informative than a' if

$$\psi_2^d(a) \geq \psi_2^d(a'). \tag{38}$$

Suppose $\psi_1^s(\cdot)$, $\psi_1^d(\cdot) > 0$. Then to ensure upward sloping supply we require that $\psi_2'^s(\cdot) > 0$ and for the demand curve to be downward sloping we require that $\psi_2'^d(\cdot) < 0$.[15] With this framework, the bid–ask spread is narrower under the active learning strategy. Intuitively, price discovery takes the form of inducing a higher signal-to-noise ratio by including a larger systematic order flow relative to the volume of noise trading. As $\psi'^d < 0$ and $\psi'^s > 0$, it follows that price discovery requires lower ask prices and higher bid prices, i.e., a smaller spread. Formally, suppose the single-period maximizing prices are $a_0^M \geq b_0^M$. Again, our assumptions concerning the independence of ε_t^s and ε_t^d along with the common valuation target θ imply that we can separate the two sides of the problem when considering learning. For any fixed bid price, say $\bar{b}$, suppose $a_0^* > a_0^M$. As $\psi_2'^d(\cdot) < 0$, a_0^M is more informative than a_0^*. Therefore, a_0^M is preferred to a_0^*, which contradicts the assumption that a_0^* is the optimal decision. So, $a_0^* \leq a_0^M$. A similar argument shows that $b_0^* \geq b_0^M$. So, if $a_0^* > a_0^M$ or $b_0^* <$

[15] We have chosen this particular configuration for the functional specification so that the ψ_2 functions incorporate the normal slope characteristics associated with supply and demand functions. It should be noted, however, that placing part of the slope characteristics in the ψ_1 functions can alter and even reverse the bid–ask spread implications. For example, if $\psi_1^d < 0$ and $\psi_1^s < 0$ then for downward-sloping demand and upward-sloping supply, we must have $\psi_2'^d > 0$ and $\psi_2'^s < 0$ indicating that the slopes of demand and supply curves are opposite in sign to the slopes of the ψ_2's. It can easily be shown that, in contrast to our example, this specification results in experimental spreads which are wider than myopic spreads.

b_0^M, the market maker can improve the current period expected profit by setting the inequality in question to an equality. Further, since this action is more informative, the value function would increase with this action. Therefore, to avoid contradiction, $a_0^* \leq a_0^M$ and $b_0^* \geq b_0^M$. This example shows that experimentation can induce a narrower bid–ask spread than the single-period optimal choice.

We had noted earlier that event studies based on returns estimated around information events may be misleading if market makers adopt active learning strategies to facilitate price discovery. This example demonstrates that similar concerns apply to empirical studies using changes in the size of the bid–ask spread as evidence of information-based trading. This result is, in part, a direct consequence of a transfer of trading profit from the market maker to early traders in return for more accurate information regarding the fundamental value of the security.[16]

3.4. *Example 4: A Numerical Example*

In order to illustrate that the inequality in Proposition 2 can be strict, we used numerical methods to compute the optimal and myopic strategies for a linear model with unknown slope coefficient. Specifically, we have $\psi_1(\beta) = -\beta$, $\psi_2(p) = p$, and $\beta \in \{\beta^1, \beta^2\}$. The initial parameter values are as follows: $\mu_0 = 0.5$, $T = 2$, and $\{\beta^1, \beta^2\} = \{20, 30\}$. The quantities are interpreted as round lots of 100 shares. We allowed the intercept of the excess demand function, ψ_0, to vary between 100 and 1000, in increments of 20. The full information market-clearing price in this model is $\theta = -\psi_0/\psi_1$, so that the parameter values imply (expected) market-clearing prices ranging from 3.33 to 50. In order to assess the impact of noise, we also allowed the standard deviation of the error term, σ, to vary from 20 to 1000, in increments of 20. For computational reasons, we chose the price set I to be discrete in increments of one cent.

Figure 1 displays the absolute deviation from the myopic price. By Proposition 2, this deviation is nonnegative. For the displayed range of parameters, the optimal price reached a maximum of 1.375% above the single-period price at $(\psi_0, \sigma) = (200, 40)$ and the largest absolute deviation was 53 cents at $(\psi_0, \sigma) = (1000, 180)$ and at $(1000, 200)$. Extending the time horizon should increase the benefit from experimentation (e.g., Prescott, 1972), and hence the "abnormal" return following an information event.

Perhaps most interestingly, the figure demonstrates that the relationship of experimentation-induced price deviation to noise is not monotone. With low levels of noise, there is no benefit to experimentation because

[16] Gammill (1990) makes a similar point, but in the context where a market maker knowingly trades with an insider to infer her private information. Here, we have a noncooperative side-payment, in expectation, to informed traders in return for more accurate information.

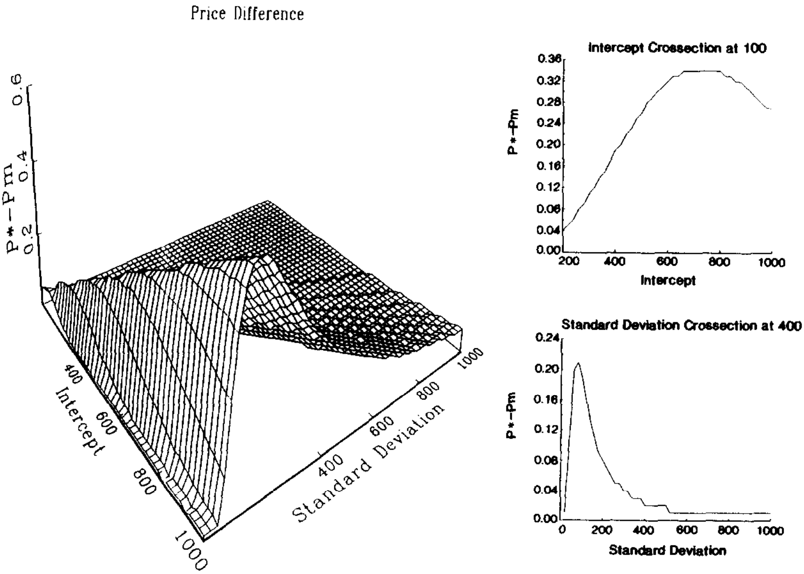

FIG. 1. Divergence of optimal prices from myopic levels.

the myopic action reveals almost as much information as the optimal action. This suggests that the amount of experimentation may increase with the level of noise, at least up to a point. Beyond this point, there is little hope in learning anything from trade and current losses outweigh the expected gains from learning. There are regions in which experimentation is more likely to be observed. Put differently, stocks with intermediate noise levels (where the meaning of "intermediate" depends on the fixed intercept) are better candidates for price experimentation than stocks with either high or low levels of noise.

The level of experimentation-induced deviation from myopia also is nonmonotone in the intercept term. A greater intercept implies a larger unconditional variance of the market-clearing price as well as a higher price; for a given standard deviation, the spread between the two possible clearing prices is $\psi_0(1/\beta^1 - 1/\beta^2)$, which is linear in ψ_0. Larger intercepts also increase the asset's expected value, and the combination of both factors tends to increase the incentive to experiment. However, this effect can be offset because larger intercepts also imply that any given *percentage* deviation from the expected value of the asset involves a larger *absolute* expected cost. The same is true for the myopically optimal price. Hence, as the intercept grows large, p^M and p^* converge. So, as was the case for noise, stocks with intermediate levels of known price-insensitive excess demand (where the meaning of "intermediate" depends on the fixed level of noise) are more likely to be candidates for price experimentation.

To summarize, the relation between price uncertainty, demand uncertainty, and the incentives to engage in costly price discovery is more complex than appears at first glance. Our example suggests that experimentation may influence prices and trade-to-trade returns. The magnitude of the effect varies over the levels of the two types of price-insensitive demand in our model: noise and known price-insensitive demand (the intercept).

4. IMPLICATIONS OF THE PRICE DISCOVERY HYPOTHESIS

In asymmetric information models such as Kyle (1985) or Glosten and Milgrom (1985), price equals the expected value of the security and the expected net order flow following a revision in the market maker's quoted prices is zero. In our model, however, changes in price quotations are associated with nonzero expected order flow because the market maker's price discovery strategy is to raise the signal-to-noise ratio by deviating from the single-period optimal price. Thus, price discovery implies causality running from quote revisions to order flow. Evidence consistent with this prediction is reported by Hasbrouck (1991), who estimates vector autoregressions using intraday data for NYSE stocks, and finds evidence that quote revisions cause, in the sense of Granger–Sims, trades. There is an inverse relation between quote revisions and subsequent trades. Hasbrouck notes that this finding is consistent with the market makers setting quotes to optimally extract information from traders, as we suggest, but that this finding could also reflect the effect of inventory control by the specialist.[17]

Madhavan and Smidt (1992) provide evidence to help distinguish these two competing hypotheses. Using specialist data for 16 NYSE stocks in 1987 to control for inventory effects, they find evidence to suggest that the specialist's quotes anticipate future order flow imbalances in excess of those predicted on the basis of current trading information. Further, Granger–Sims causality runs from quote revisions to future inventory movements. Both of these findings are consistent with the expected order flow following a price revision being nonzero as predicted by our model, but not by existing asymmetric information models.

The price discovery hypothesis also suggests systematic deviations from the normal pattern of pricing prior to known information events. We

[17] Intuitively, a risk-averse market maker will adjust prices to control fluctuations in inventory.

know of no study that has attempted to test this hypothesis. However, Handa (1991) attempts to isolate the effect of price discovery by examining the distribution of prices conditional upon past price movements. In the prototypical asymmetric information models, such as Glosten and Milgrom (1985), the transaction-to-transaction price change reflects changes in the specialist's expectation of the asset's value. Since these price changes are innovations, there should be no predictability conditional upon past price movements. This is not necessarily true in a model of the type developed here, as Example 2 makes clear.

Handa uses intraday transaction data for roughly 2000 NYSE and AMEX stocks for 1988 and 1989 to examine the future distribution of price changes conditional upon the history of quote revisions. The basic idea is that large quote revisions signal future volatility, and hence are likely to be associated with attempts to expedite price discovery by market makers. Handa finds that the conditional distribution of prices violates the martingale property conditional upon a price movement, which is inconsistent with the standard asymmetric information models but is consistent with our model. Handa concludes that the "price discovery–price experimentation hypothesis provides the most consistent explanation of the evidence" that he documents. It would be interesting to repeat these tests using intraday transactions data drawn from periods preceding known information events, since this is when price discovery is most likely to occur. This is a topic for further research.

Our analysis has several policy implications for the design of trading protocols and structures. First, a common suggestion to reduce price volatility is a trading halt triggered by a large price movement or an order imbalance. Our analysis suggests "abnormal" price movements and order imbalanced may reflect price discovery by market makers which helps the market to resolve uncertainty more rapidly. Thus, a security-specific trading halt may exacerbate the problem it is designed to remedy. On the other hand, after a halt most continuous markets reopen trading with a call (batch) auction that could provide more efficient price discovery than a continuous mechanism.[18] Second, our model suggests that competition will not generate sufficient investment in the production of information because free entry undermines the incentives to follow active learning strategies and thereby collect costly information. This suggests a reason for the dominance of specialist systems over time. Price discovery has aspects of a public good, which is not widely recognized in policy discussions. We note that price discovery can occur even in markets with little uncertainty concerning asset values. From a policy perspective, concern about price discovery need not be confined to thinly traded assets where asymmetric information is likely to be a problem.

[18] See also Madhavan (1992) for a discussion of this point.

In general, the incentives for market makers to undertake costly price discovery depend on market microstructure; previous discussions of the relative merits of various market structures have not recognized price discovery as an issue. Our analysis suggests that price discovery may be a consideration in the evaluation of market mechanisms.

The discussion above, however, must be evaluated in the light of two potentially serious limitations of the analysis. First, the assumed trading protocol is simplified for tractability, and the correspondence to real-world trading structures is weak. In particular, there may be some low cost methods of providing for price discovery in real-world markets. Indeed, most market mechanisms provide methods for market participants to "test the waters at low cost." For example, specialists on the NYSE can limit their losses by restricting market depth. This would tend to increase the incentives to use prices as a tool for expediting price discovery. Conversely, many batch markets allow for costless experimentation by floor traders in the form of price indications which are nonbinding. Such opportunities may alter the incentives to undertake price discovery through price experimentation. Second, the assumption that the informed traders do not attempt to strategically manipulate the market maker's beliefs is restrictive, especially if there are relative few informed traders. One rationale for our approach is that informed traders may face a "prisoners' dilemma" problem in that attempts to game the market maker by one insider is open to cheating by other insiders.

Despite these limitations, the model's basic conclusions appear to be robust. Market makers have the incentives and the ability to undertake costly price discovery through a process of testing the waters. Models that ignore such behavior may be overlooking an important dimension of dynamic trading strategies. Whereas our model may be difficult to test directly, it appears to be consistent with stylized facts that are difficult to explain using existing paradigms.

5. CONCLUSIONS

We demonstrate that the market maker has the ability and incentive to perform the critical function of price discovery. We model the decision problem underlying this function and show that dynamic pricing may involve departures from the single-period optimal choice. By doing so, the market maker induces order flows that are more informative in a statistical sense than the order flow at the single-period price. Although these strategies are costly in the short run, they are optimal in an intertemporal framework because better current information increases future expected utility. Intertemporal price discovery by a market maker takes the form of price experimentation.

We examine four special cases of a general model and find that price discovery is not optimal for certain functional forms of demand. When traders have negative exponential expected utility functions and the future value of the asset is normally distributed, market makers have no incentive to engage in price discovery, even though trading reveals traders' private information. However, we provide an example where price discovery occurs even though all agents (including the market maker) have homogeneous expectations regarding the full information value of the asset. In this case, price discovery occurs because it is in the market maker's interest to learn about traders' elasticities of demand. We also show that price discovery can affect the size of the bid–ask spread. With a numerical example, we show that price discovery may affect returns, and that price discovery is related in a nonmonotonic fashion to uncertainty regarding fundamentals and to the level of liquidity-motivated trading. The numerical example suggests that returns preceding an information event may reflect departures from the standard pricing model. These abnormalities associated with price discovery may bias event studies.

To summarize, the paper's main contribution is to demonstrate that market makers play an important role in expediting price discovery beyond that of simply providing liquidity. While price discovery is costly to market makers, we showed that under certain conditions they have incentives to bear these costs. Our examples illustrate the complexity of this activity; while the idea of price discovery through experimentation is quite intuitive, the examples show that this intuition can be misleading. The model has implications for empirical studies of security prices and for policy proposals regarding security market microstructure.

APPENDIX

Proof of Lemma 1. Our approach follows GKM with certain important modifications necessitated because the unknown parameter, β, enters the utility (expected profit) function. The proof omits some technical details for expositional ease. We begin by constructing the decision set of the agent. For $t \geq 1$, let γ_t denote a function with domain H_{t-1} and range I, and let $\Gamma_t = \{\gamma_t: H_{t-1} \rightarrow I\}$. Define $\gamma_T = \{\gamma_t\}$, and $\Gamma_T = \Gamma_1 \times \cdots \Gamma_T$. The interpretation is that γ_T is a vector of contingency plans, one for each period which is measurable with respect to the market maker's information set. That is, they are functions of the observable history. Define $T + 1$ functionals ζ_t: $\Gamma_T \times \Re^t \times I \rightarrow I$ by

$$\zeta_0(\gamma_T, p_0) = p_0$$

$$\zeta_1(\gamma_T, z_0, p_0) = \gamma_1(z_0, p_0)$$

$$\vdots$$

$$\zeta_t(\gamma_T, z_{t-1}, p_{t-1}) = \gamma_t(z_{t-1}, \zeta_{t-1}(\gamma_T, z_{t-2}, p_{t-2}))),$$

where $\zeta_{t-1} = \zeta^k$ for $k = 0, \ldots, t-1$. The interpretation of these functionals is that each realization of the contingency plan (except for the first) is selected as a function of the previous choices (p_t's). Therefore, p_t is actually dependent on previous excess demands (z_t's) and the original price choice (p_0). The relationship of the decision to these variables is specified by the function γ. The rule ζ_t ties the decision p_t to the history of excess demands, the initial price and the contingency plan vector. The prices chosen, for any observable history, are $p_t = \zeta_t(\cdot, \cdot, \cdot)$. The market maker's dynamic optimization problem is to choose an initial price, $p_0 \in I$, and the T-vector of functions $\gamma_T \in \Gamma_T$ to maximize the discounted expected profit. This completes our construction of the action set.

Next, we define the function η_T: $I \times \Gamma_T \times \beta \rightarrow \Re$ by

$$\eta_T(p_0, \gamma; \beta^i) = \left[\phi(p_0; \beta^i) + \sum_{t=1}^{T} \delta^t \int_{\Re_t} \phi(p_t; \beta^i) \prod_{\tau=0}^{t-1} h(z_\tau; f(p_\tau; \beta^i)\, dz_{t-1}\right], \tag{39}$$

where the function $\phi(p; \beta)$ is the expected utility function conditional upon β:

$$\phi(p; \beta^i) \equiv \int_{\Re} u(p, \beta^i, f(p, \beta^i) + \varepsilon) h(\varepsilon; 0)\, d\varepsilon. \tag{40}$$

Then, we can show that

$$V_T(\mu_0) = \sup_{p_0, \gamma} \sum_i \eta_T(p_0, \gamma; \beta^i) \mu_0^i. \tag{41}$$

The formal derivation of this equation is omitted here because of its length and the complexity added by state-dependent utility, but it is straightforward using the definition of $V_T(\mu_0)$. Heuristically, the $\gamma_t(\cdot)$ functions specify optimal actions from the set $I \times \Gamma_T$ for all possible states of the world, which is exactly equivalent to the maximization problem (5) where the actions were prices, concluding the proof. ■

Proof of Lemma 2. See DeGroot (1970), p. 436).

Proof of Lemma 3. Necessity is obvious. For sufficiency, we proceed inductively. Given that

$$\boldsymbol{\mu}_1(z_0, p_0, \mu_0) = \boldsymbol{\mu}_1(z_0', p_0', \mu_0), \tag{42}$$

the induction hypothesis implies

$$\boldsymbol{\mu}_s(\mathbf{z}_{s-1}, \mathbf{p}_{s-1}, \mu_0) = \boldsymbol{\mu}_s(\mathbf{z}_{s-1}', \mathbf{p}_{s-1}', \mu_0) \tag{43}$$

for all $s = 1, \ldots, t$, where $\mathbf{z}_{s-1}'$, and $\mathbf{p}_{s-1}'$ are the vectors of past realizations of excess demand and prices where the first price is p_0' rather than p_0. Since actions depend on only the posterior and time, it follows that $p_t = p_t'$ and $z_t = z_t'$ for $t \geq 1$. Applying Bayes' rule, we obtain

$$\mu_{t+1}^i = \frac{\mu_t^i h(z_t - f(\beta^i, p_t(\mu_t)); 0)}{\sum_{k=1}^n \mu_t^k h(z_t - f(\beta^k, p_t(\mu_t)); 0)}, \tag{44}$$

By the induction hypothesis and the fact that the realizations at time t (z_t's) are the same, we can write

$$\mu_{t+1}^i = \frac{\mu_t^{i'} h(z_t' - f(\beta^i, p_t(\mu_t')); 0)}{\sum_{k=1}^n \mu_t^{k'} h(z_t' - f(\beta^k, p_t(\mu_t')); 0)}, \tag{45}$$

which is just $\mu_{t+1}^{i'}$. ∎

Proof of Proposition 1. We must show that functional forms satisfying the derivative condition (and only those) dictate posterior beliefs which are invariant under current price choice. By Lemma 1, invariance of one-period forward posteriors in price choice is all we must consider.

(*Necessity*) Let τ denote the true state. Then invariance in the one-period forward posterior requires

$$\frac{\mu_0^i h(f(p, \beta^\tau) + \varepsilon - f(p, \beta^i); 0)}{\sum_{k=1}^n \mu_0^k h(f(p, \beta^\tau) + \varepsilon - f(p, \beta^k); 0)} = \frac{\mu_0^i h(f(p', \beta^\tau) + \varepsilon - f(p', \beta^i); 0)}{\sum_{k=1}^n \mu_0^k h(f(p', \beta^\tau) + \varepsilon - f(p', \beta^k); 0)} \tag{46}$$

for all i; for all $p, p' \in I$. It follows that

$$\frac{h(f(p, \beta^\tau) + \varepsilon - f(p, \beta^i); 0)}{h(f(p', \beta^\tau) + \varepsilon - f(p', \beta^i); 0)} = \frac{h(f(p, \beta^\tau) + \varepsilon - f(p, \beta^j); 0)}{h(f - (p', \beta^\tau) + \varepsilon - f(p', \beta^j); 0)} \tag{47}$$

for all i, j; for all $p, p' \in I$. Set $i = \tau$ on the left-hand side to get

$$h(f(p, \beta^\tau) + \varepsilon - f(p, \beta^j); 0) = h(f(p', \beta^\tau) + \varepsilon - f(p', \beta^j); 0). \quad (48)$$

This holds for all ε and therefore

$$f(p, \beta^\tau) + \varepsilon - f(p, \beta^j) = f(p', \beta^\tau) + \varepsilon - f(p', \beta^j), \quad (49)$$

which implies that for all $p, p' \in I, f(p, \beta^\tau) - f(p, \beta^j)$ does not depend on p. It follows then that $f(p, \beta^i) - f(p, \beta^j)$ does not depend on p and that $f(p, \beta^i) - f(p', \beta^i) = f(p, \beta^j) - f(p', \beta^j)$ which directly implies that the derivative of f with respect to p does not depend on β.

(*Sufficiency of the condition for invariance*) If $\partial f/\partial p$ is independent of β, it follows that

$$f(p, \beta^i) - f(p, \beta^j) = f(p', \beta^i) - f(p'; \beta^j) = k(\beta^i, \beta^j) \quad (50)$$

for all $p, p' \in I$; for all $i, j = 1, \ldots, n$. Then

$$h(f(p, \beta^\tau) + \varepsilon - f(p, \beta^i); 0) = h(k(\beta^\tau, \beta^i))$$
$$h(f(p'; \beta^\tau) + \varepsilon - f(p', \beta^i); 0) = h(k(\beta^\tau, \beta^i))$$

for all i. This implies equal numerators and denominators in the Bayes formulae and thus equal posteriors completing the proof of sufficiency.

If controlled learning (experimentation) is futile, incurring opportunity losses by setting $p_t \neq p_t^M$ is suboptimal, and the result follows. ■

Proof of Lemma 4. Consider the experiments (see Definition 1 above) $h(z; p, \beta)$ and $h(z'; p', \beta)$. Construct the random variables z and z' to have a bivariate normal distribution with means $\psi_0 + \psi_1(\beta)\psi_2(p)$, and $\psi_0 + \psi_1(\beta)\psi_2(p')$, respectively, and variance σ^2, and set the correlation between z and z' to be $\rho = \psi_2(p')/\psi_2(p)$. Since $\psi_2(p) > \psi_2(p')$, $p < 1$ as required. (Remember $\psi_2 > 0$.) Using the projection theorem for the normal distribution, the conditional mean of z' given z is

$$E[z'|z] = E[z'] + \rho(z - E[z]) \quad (51)$$

$$= \psi_0 + \psi_1(\beta)\psi_2(p') + \frac{\psi_2(p')}{\psi_2(p)}(z - \psi_0 - \psi_1(\beta)\psi_2(p)) \quad (52)$$

$$= \psi_0\left(1 - \frac{\psi_2(p')}{\psi_2(p)}\right) + \frac{\psi_2(p')}{\psi_2(p')z}, \quad (53)$$

which is independent of β. Similarly, the conditional variance of z' given z is

$$\sigma^2[z'|z] = \sigma^2\left(1 - \left[\frac{\psi_2(p')}{\psi_2(p)}\right]^2\right), \quad (54)$$

which is positive since $\psi_2(p) > \psi_2(p') > 0$. Denote by $\nu(z'; z)$ the conditional distribution of z' given z, which we know to be normal. Now, by construction $h(\cdot; \cdot)$ is the marginal density of z and $h'(\cdot; \cdot)$ is the marginal density of z'. We know that the marginal density of z' is related to the marginal density of z and the conditional density of z' given z as

$$h(z'; \cdot) = \int_{\Re} \nu(z'; z) h(z; \cdot)\, dz, \tag{55}$$

where we know $\int_{\Re} \nu(z'; z)\, dz' = 1$. Comparing (55) and Definition 1, we see that the experiment $h(z; p, \beta)$ is more informative than the experiment $h(z'; p', \beta)$, because $\nu(z'; z)$ is independent of β. ■

Proof of Proposition 2. In order to produce a contradiction suppose $\psi_2^* < \psi_2^M$. By the uniqueness of the myopic policy, it must be that

$$U(p_0^M; \mu_0) > U(p_0^*; \mu_0) \tag{56}$$

for $p_0^M \in P_0^M$, and $p_0^* \in P_0^M$. From Lemma 4, the experiment with $\psi_2(p^M)$ is more informative than the experiment with $\psi_2(p^*)$. Then, Blackwell's Theorem implies

$$E[V_{T-1}(g_1(\mu_0; p_0^M, z_1^M))] \geq E[V_{T-1}(g_1(\mu_0; p_0^*, z_1^*))]. \tag{57}$$

Therefore

$$U(p_0^M; \mu_0) + \delta E[V_{T-1}(g_1(p_0^M; \mu))] > U(p_0^*; \mu_0) + \delta E[V_{T-1}(g_1(p_0^*; \mu))], \tag{58}$$

which contradicts the definition of ψ_2^*. Therefore $\psi_2^* \geq \psi_2^M$. ■

Proof of Proposition 3. Using the fact that $\tilde{\varepsilon}_t$ is uncorrelated with $\tilde{\beta}$, the variance of the initial change in market maker inventories is given by

$$\mathrm{Var}(\tilde{z}_0) = \psi_2(p)^2 \,\mathrm{Var}[\psi_1(\tilde{\beta})] + \sigma^2.$$

Then, the first part follows directly from Proposition 2. With $u(p, \beta, z) = (p - F^{-1}(0))z$, where $F(\cdot) = f(\cdot; \beta)$, active experimentation implies that $U(p_0^M; \mu_0) \geq U(p_0^*; \mu_0)$, which establishes the second statement. ■

REFERENCES

ADMATI, A., AND PFLEIDERER, P. (1988). A theory of intraday trading patterns, *Rev. Finan. Stud.* **1,** 3–40.

BERRY, D., AND FRISTEDT, B. (1985). "Bandit Problems: Sequential Allocation of Experiments." Chapman & Hall, London.

BLACKWELL, D. (1953). Equivalent comparisons of experiments, *Ann. Math. Statist.* **24,** 265–272.

BRONFMAN, C., AND SCHWARTZ, R. (1990). "Trading, Noise, and Market Structure," Working paper, New York University.

DEGROOT, M. (1970). "Optimal Statistical Decisions." McGraw–Hill, New York.

DEMSETZ, H. (1968). The cost of transacting, *Quart. J. Econ.* **82,** 33–53.

EASLEY, D., AND KIEFER, N. (1988). Controlling a stochastic process with unknown parameters, *Econometrica* **56,** 1045–1064.

EASLEY, D., AND O'HARA M. (1987). Price, quantity and information in securities markets, *J. Finan. Econ.* **19,** 69–90.

GAMMILL, J. (1990). "The Organization of Financial Markets: Competitive versus Cooperative Market Mechanisms," Working paper, Harvard Business School.

GLOSTEN, L., AND MILGROM, P. (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents, *J. Finan. Econ.* **14,** 71–100.

GROSSMAN, S. (1976). On the efficiency of competitive stock markets when traders have diverse information, *J. Finance* **31,** 573–585.

GROSSMAN, S., KIHLSTROM, R., AND MIRMAN, L. (1977). A Bayesian approach to the production of information and learning by doing, *Rev. Econ. Stud.* **44,** 533–547.

GROSSMAN, S., AND MILLER, M. (1988). Liquidity and market structure, *J. Finance* **42,** 617–637.

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

HANDA, P., AND SCHWARTZ, R. A. (1989). "Dynamic Price Discovery and the Short-Term Behavior of Security Prices," Working paper, New York University.

HASBROUCK, J. (1991). Measuring the information content of stock trades, *J. Finance* **46,** 178–208.

HO, T., AND STOLL, H. (1983). The dynamics of dealer markets under competition, *J. Finance* **38,** 1053–1074.

KIHLSTROM, MIRMAN, R. L., AND POSTELWAITE, A. (1984). Experimental consumption and the 'Rothschild Effect' *in* "Bayesian Models in Economic Theory" (M. Boyer and R. Kihlstrom, Eds.), pp. 279–302. Elsevier, Amsterdam/New York.

KYLE, A. (1985). Continuous auctions and insider trading, *Econometrica* **53,** 1315–1335.

LEACH, J. C., AND MADHAVAN, A. (1991). "Price Experimentation and Security Market Structure," Working paper, University of Pennsylvania.

MADHAVAN, A. (1992). Trading mechanisms in securities markets, *J. Finance* **47,** 607–642.

MADHAVAN, A., AND SMIDT, S. (1992). "An Analysis of Daily Changes in Specialist Inventories and Quotations," Working paper, University of Pennsylvania.

MCLENNAN, A. (1984). Price dispersion and incomplete learning in the long run, *J. Econ. Dynam.* **7,** 331–347.

O'HARA, M., AND OLDFIELD, G. (1986). The microeconomics of market making, *J. Finan. Quant. Anal.* **21,** 361–376.

PRESCOTT, E. (1972). The multiperiod control problem under uncertainty, *Econometrica* **40,** 1043–1048.

ROTHSCHILD, M. (1976). A two-armed bandit theory of market pricing, *J. Econ. Theory* **9,** 185–202.

SCHREIBER, P., AND SCHWARTZ, R. (1986). Price discovery in securities markets, *J. Portfol. Manage.* **12,** 43–48.