THE JOURNAL OF FINANCE • VOL. XL, NO. 1 • MARCH 1985

On the Feasibility of Automated Market Making by a Programmed Specialist

NILS H. HAKANSSON, AVRAHAM BEJA, and JIVENDRA KALE*

ABSTRACT

Securities trading is accomplished through the execution of orders. Admissible orders (e.g., market orders, limit orders) give rise to discontinuous aggregate demand functions, composed of many "steps." *Demand smoothing*, or the balancing of excesses due to such discontinuities via intervention, is one of the most basic functions that could be assigned to a "specialist." When the specialist's "affirmative obligation" is fully specified, his or her activity can in principle be automated. This paper is an attempt to assess, via simulation, some of the ramifications of using a "programmed specialist," whose automated market making is limited to demand smoothing. A number of alternative rules of operation are simulated. Several of the rules performed well, especially the extremely simple rule that calls for the (computerized) specialist to minimize new absolute share holdings in each security at each trading point via "total" (as opposed to "local") demand smoothing. Our results indicate that the underlying costs of demand smoothing are on the order of a fraction of a penny per share traded even in relatively thin markets.

THE FUNCTION OF THE specialist as a "designated market maker" has traditionally been taken for granted by market participants while being practically ignored by financial and economic theory. Recently, this function has become the subject of increasing scrutiny. Current discussions in connection with the possible reorganization of the exchanges into a National Securities Market have raised serious questions concerning the roles (if any) that should be assigned to specialists in this market, especially if the order clearing mechanism is to be highly automated.[1] On the theoretical side, the specialist's function has become one of the focal points in an emerging new field of financial theory which concentrates

* University of California, Berkeley, Tel Aviv University, and BARRA, Berkeley, respectively. This is an abridged version of [19], which is available upon request from the authors (School of Business Administration, University of California, Berkeley, Califronia 94720). Among other things, the comprehensive version contains a detailed description of the simulated demand generation process and a fuller discussion of the results. Presented at The Institute of Management Science Meeting in Honolulu; Simon Fraser University, Burnaby, British Columbia; Bell Laboratories, Murray Hill, New Jersey; New York University; and Tel Aviv University. The authors would like to thank the participants of these seminars for helpful comments and gratefully acknowledge support from the National Science Foundation Grant SOC77-09482, from Ford Foundation Grant No. 10, administered through the Israel Foundation Trustees, and from the Dean Witter Foundation. Portions of the model described in this paper have been adopted by the Tel Aviv Stock Exchange.

[1] A fairly comprehensive presentation of the various viewpoints is available in Bloch and Schwartz [7].

on modeling the microstructure of securities market systems.[2] At this early stage in the field's development, the microstructure literature offers no *comprehensive* theory (or theories) of the specialist's role. Rather, individual studies relate to *specific* aspects of the specialist's activity.[3] Collectively, these studies attempt to identify circumstances in which intervention in the process of order execution is either necessary or desirable, and to evaluate alternative ways of effecting such intervention. The present paper extends this analysis and its potential contribution to policy decisions by studying in some detail the central function of a designated market-making entity. We demonstrate that the rendered service is indeed needed, that it can be offered at a very reasonable cost, and that it is also highly amenable to full automation.

In economic theory, the traditional view of the market as a Walrasian auction assumes all the necessary conditions for the existence of "equilibrium prices," i.e., prices at which the "competitive demand schedules" based on the traders' underlying preferences are simultaneously cleared in all markets. The trades which are executed at those prices produce Pareto-efficient allocations, and in this context there is thus no need and no room for a "specialist" that trades on the basis of motives other than his or her own current portfolio preferences.

Of course, the modern theory of securities markets operation is less concerned with the existence of "ideal" trades than it is with the way trades can operationally be attained. In practice, trades must be based on outstanding *orders*, which form an unavoidable link between the traders with their underlying preferences and the market system (this is discussed in considerable detail in Beja and Hakansson [4]). One of the critical limitations on the efficient operation of actual securities markets thus involves the inherent differences that must exist between the orders submitted in those markets and the idealized demand schedules of economic theory. For example, today's exchange procedures do not admit "joint" limit orders that condition transactions in one asset on the prices of other assets. In addition, a trader cannot, as a practical matter, submit even a rough approximation of a continuous demand schedule in the great majority of cases but must be content with the opportunity to submit piecewise linear segments (in the form of "market" and "limit" orders). Consequently, the natural trades that would be generated by the current *orders* may differ from the ones that would be most appropriate for the current profile of the investors' true *preferences*. More importantly, the discrete nature of the submitted orders gives rise to a *discontinuous aggregate excess demand function*, so that a price that clears the outstanding orders *will generally fail to exist.*[4] Thus, for trading to occur, the use of discrete

[2] Microstructure theory is devoted to explicit study of such aspects as the trade execution mechanism and the behavior of the participating agents and institutions. For a very limited sample, see Garman [17], Beja and Hakansson [4], Cohen, Maier, Schwartz, and Whitcomb [9–11], Garbade and Silber [16], Stoll [27], Goldman and Beja [18], Beja and Goldman [3], Garbade and Sekaran [15], Ho, Schwartz, and Whitcomb [22], and Mendelson [24, 25].

[3] For earlier work in this area, see Black [6], Smidt [26], Tinic and West [28], and Barnea [2]. For more recent studies, see, e.g., Beja and Hakansson [5], Goldman and Beja [18], and Cohen, Maier, Schwartz, and Whitcomb [13, 14]. See also Garman [17], Amihud and Mendelson [1], and Mendelson [24, 25].

[4] The easiest way to see this is to consider two discrete functions, one upward sloping and one downward sloping; two such functions generally have no point of intersection.

orders inevitably requires either 1) direct intervention by a "market maker," or 2) the rationing of either buy or sell orders. Having a specialist who trades for his or her own account is the most convenient way to implement the market-making approach.[5] We call the function of balancing excesses due to discontinuities of the aggregate orders *demand smoothing.* "Keeping the market going" by carrying out this function is the affirmative obligation of the specialist as a designated market maker.[6]

To illustrate, suppose that the most recent price of a given security is \$10 per share, and that the currently outstanding orders are: (1) a market sell order for 200 shares, (2) a limit sell order for 200 shares at \$9½ or better, and (3–4) two limit buy orders, for 300 shares each, at \$9¾ or less. Then, at all prices below \$9½ there is an excess demand of 400 shares, at all prices between \$9½ and \$9¾ inclusive there is an excess demand of 200 shares, and at all prices above \$9¾ there is an excess supply of 400 shares, so that a market clearing price does not exist (because the traders' aggregate submitted demand is a discontinuous function of the price). Clearly, the above orders could be entries in the "book" in either a continuous market or a call auction environment.

In the preceding example, there are numerous alternatives available to the market-making specialist in clearing the market by "smoothing" the discontinuous demands submitted by the traders: he or she can buy 400 shares at some price above \$9¾, sell 200 shares at a price between \$9½ and \$9¾ inclusive, or sell 400 shares at some price below \$9½. Or a priority-oriented rule (of the type used on the New York Stock Exchange), with queues based on price, time, size, etc., or a rule based on rationing can be used. Which choice is best? It is the

One might, as a practical matter, wonder whether discontinuities in outstanding demand are really typical phenomena that significantly hamper the trading process. It might superficially seem, for example, that the aggregation of a large number of orders, where the few orders of each investor constitute a substantially discontinuous approximation to underlying demand schedule, would tend to become approximately continuous as the number of traders gets very large. However, a number of arguments indicate that such discontinuities are not untypical even in rather thick markets. First, note that trading in the major exchanges in the U.S. is for the most part limited to prices which are set in exact eighths of a dollar, so that any down-sloping demand must have "discontinuities" (these discontinuities would not be alleviated by going to "decimal" pricing, for example). Second, the price limits that individual investors choose to set for their limit orders are not purely random choices. Rather, they are all jointly related to the previous price, and would in many cases tend to cluster around a relatively small number of "natural" choices. (See Cohen, Maier, Schwartz, and Whitcomb [12] for a model where traders use elaborate strategies in determining the price limit in their orders.) Ultimately, the strongest argument for the empirical relevance of discontinuities in the investors' aggregate submitted demand is the persistence of substantial bid-ask spreads in security prices; if aggregate demand were practically continuous, no bid-ask spreads would remain effective.

[5] Due to space limitations, rationing is not considered in this paper. For an analysis of such approaches to securities trading, see Beja and Hakansson [4, 5, pp. 148–52].

[6] Other dimensions may, of course, be imputed to the specialist's "obligation." For example, it may be desirable for the specialist to intervene to avoid "extreme" price fluctuations caused by small orders if it is believed that the orders reflect some idiosyncratic random phenomenon which is unrelated to basic values (as reflected by investors' preferences and beliefs). In other words, the specialist can intervene in order to offset *thinness in outstanding orders* and preserve the current price if he or she believes that this price reflects best the investors' overall underlying demands. See also Goldman and Beja [18].

purpose of this paper to throw light on the implications of some of the possible decision rules that might be followed.[7]

In the present paper, we are interested in the "demand smoothing" aspect of the market-making functions of a specialist operating in an organized exchange.[8] In particular, we attempt to assess, via simulation, some of the ramifications of using a "fully automated specialist" whose task is limited to "demand smoothing." The analysis is restricted to rules under which all orders received in a given period are executed at the same price and in full.[9,10] Within this family, a number of alternative, and initially equally plausible, rules of operation are simulated. As described in more detail in Section II, all of the rules examined seek to in some sense minimize the "noise" introduced by the microstructure environment. For each rule, we evaluate the extent of the specialist's participation in trading, the behavior of stock and cash positions, and net profits or losses. By implication, this is indicative of the basic costs of alternative methods of demand smoothing. In addition, we study how the preceding quantities depend on the number of securities that the specialist handles.

The paper proceeds as follows. A general description of the simulation model is presented in Section I. Section II describes the different rules of market making investigated in this study. The results of the simulations are summarized in Section III. Section IV examines the implications relative to the social costs of the demand smoothing function, while Section V contains a concluding summary.

I. The Simulation Model

Our model describes an environment with I *investors* (synonomously also called *traders*) trading in S *securities*. When an investor wishes to trade, one or more *orders* are submitted to a central marketplace where these orders are recorded in the "book" and executed according to a well-specified operating procedure. *Trading periods* form a sequence of points, which may possibly be randomly spaced on the time axis, and we index these points by $t = 1, 2, \cdots$. All trades in a security executed at the same period have the same dollar price per share, stated in eighths of a dollar. Short positions are possible. There is a single

[7] In current practice, the human specialist intervenes on the basis of "professional judgment." However, the specialist's discretion in setting prices is obviously not unlimited. It is effectively constrained by explicit rules and by mores, through external control, and through self-imposed etiquette.

[8] See also Beja and Hakansson [5, pp. 153–59].

[9] Whether the exchange is thought of as operating in "continuous" or "call auction" fashion is essentially unimportant since, from the perspective of demand smoothing, the central difference between the two is the thinness or thickness of orders at the time of execution. However, since the execution of many orders simultaneously at the same price, and an absence of priority features, are more characteristic of call markets than continuous markets, the present setting is probably somewhat closer in spirit to a call auction environment.

For a comparison between batch or call trading and continuous trading, see Ho, Schwartz, and Whitcomb [21, esp. pp. 4–5] and [22]. For a recent review of the trading system at the New York Stock Exchange, see Carrington [9].

[10] See Cohen, Maier, Schwartz, and Whitcomb [14] for a simulation study of a continuous auction market from a somewhat different perspective.

"specialist," who operates as a "designated market maker" and takes short and long positions according to some rule. His or her *portfolio* is denoted by the vector $q_t^0 = (q_{1t}^0, q_{2t}^0, \cdots, q_{St}^0)$, where q_{st}^0 is the number of shares of security s held by the specialist immediately after the trading at time t is completed. The post-trading position at time t will be summarized by the pair (c_t, v_t), where c_t is the amount of cash on hand and v_t is the value of the security portfolio. We study the behavior of (c_t, v_t) and q_t^0 on the assumption that the specialist received no income, either on capital or from his or her services, and that the cost of capital (e.g., when c_t is negative) is zero. One reason for this assumption is that one of our objectives is to ascertain what kinds of fees the specialist needs in order to perform his or her services under different market-clearing rules.

Letting $\bar{q}_t^0 = (\bar{q}_{1t}^0, \cdots, \bar{q}_{St}^0)$ denote the specialist's vector of share purchases at time t and P_t the vector of effective trading prices at time t, we have

$$\left.\begin{aligned} q_t^0 &= q_{t-1}^0 + \bar{q}_t^0, && (1)\\ v_t &= q_t^0 P_t = q_{t-1}^0 P_t + \bar{q}_t^0 P_t, && (2)\\ c_t &= c_{t-1} - \bar{q}_t^0 P_t && (3)\end{aligned}\right\} \quad t = 1, 2, \cdots$$

For each reference, we shall use the origin as our starting point, and let $c_0 = v_0 = q_{s0}^0 = 0$, for $s = 1, \cdots, S$.

The simulation model has two main parts. In the first part, which we call *the demand generation process*, we simulate the generation of the investors' orders. In each trading period, the generated orders, in aggregated form, then activate the other main part of the simulation model—the *trade execution mechanism*, which uses an automated market-making specialist to achieve market clearing.

A. The Demand Generation Process

The investors' orders are generated in each trading period by simulating their underlying (continuous) demand functions, and then approximating each demand function by a random number of orders. New demand functions are generated by jointly correlated random shifts in earlier demand functions. Each trader's submitted demand with respect to a given security may involve a market buy or sell order, one limit buy or sell order, or a set of limit buy and/or sell orders at various price limits.[11] The trader may also choose to avoid trading in a given period and submit no orders. Changes in the investors' demands for the different securities are positively correlated to reflect changes in net asset positions as well as factors common to the whole market. The shifts also involve idiosyncratic changes in each investor's relative demands for the various securities, reflecting time changes in his or her tastes or current beliefs. The demand shifts follow a nonstationary geometric random walk in which expected demands, and hence expected prices, increase over time. All demands are decreasing in price.[12] The

[11] Market and limit orders are in essence "discretized" demand schedules, which may be viewed as cost-benefit efficient attempts to capture the essential features of the idealized demand schedules. Market and simple limit orders approximate demand schedules via two vertical line segments or less, while more multiple limit orders employ three or more line segments (see [19]).

[12] The complexities introduced by relaxation of this property are not considered in this paper.

limit orders involve price limits which are restricted to a ±20% neighborhood of the previous price for the asset in question. If the submitted demands happen to drive the price in one jump beyond this specified neighborhood, the model withholds trading in the security but signals the boundary price in question to investors. As in real-world markets, individual orders for different assets are segmented (although correlated). That is, truly joint orders in which, for example, one order is conditioned on the execution of another have been ruled out as unacceptable by the market system.

Compared with the major exchanges, our simulation model appears to have a relatively small number of investors, with a relatively large average number of orders per investor and a relatively low percentage of market orders. Note, however, that the aggregation of all market orders will at most shift the aggregate demand to the left or to the right, but will not affect discontinuities so that the exact proportion of market orders is of no real consequence from the perspective at hand. In addition, patterns of aggregate demand, very similar to the ones generated in our simulation, also arise when there are many more traders submitting fewer orders. Thus, our choice of configuration for the demand generating process is in the spirit of efficient model design. As is well known, it is good professional practice in simulation studies not to attempt an "exact" replication of the environment, but rather to "amplify" intentionally the pertinent aspects.[13] Thus, we simulate an extremely thin market not because we believe it to be typical, but because it is precisely in such markets that the specialist's market-making function is most important.

B. The Trade Execution Mechanism

For each security, the investors' orders generated by the demand generation process are recorded in "the book," and summed up to give the security's aggregate demand schedule, showing for each tentative price the net outstanding orders for that security at that price. An example of an aggregate excess demand schedule, plotted for prices which are given in eighths of a dollar, is presented in Figure 1. The aggregate demand function is (necessarily) down sloping and (typically) quite "nonlinear." Since all price limits for the individual limit orders are set within a range of ±20% of the previous price, the aggregate excess demand schedule may be considered as extending vertically upwards from D_1 at the upper end (P_1) of the relevant price range, and vertically downward from D_n at the lower end (P_n) of that range. Given the discrete nature of the excess demand schedule, there is (typically) no price at which excess demand is exactly zero.

The book is governed by the simulated specialist. Given the aggregate demand schedule, the specialist selects *some* price in the relevant price range. Then, *all* the outstanding orders are executed which are effective at that price. When these orders do not happen to clear exactly, the specialist sells shares from his or her own account at the selected price if there is a positive excess demand at that price, or buys for this account if there is a negative excess demand (a net supply)

[13] For example, see Hammersley and Handscomb [20], especially the parts concerning the so-called "importance sampling." In importance sampling, one intentionally overrepresents the "more important" parts of an underlying population, so that a relevant aspect can be efficiently studied through a smaller sample.

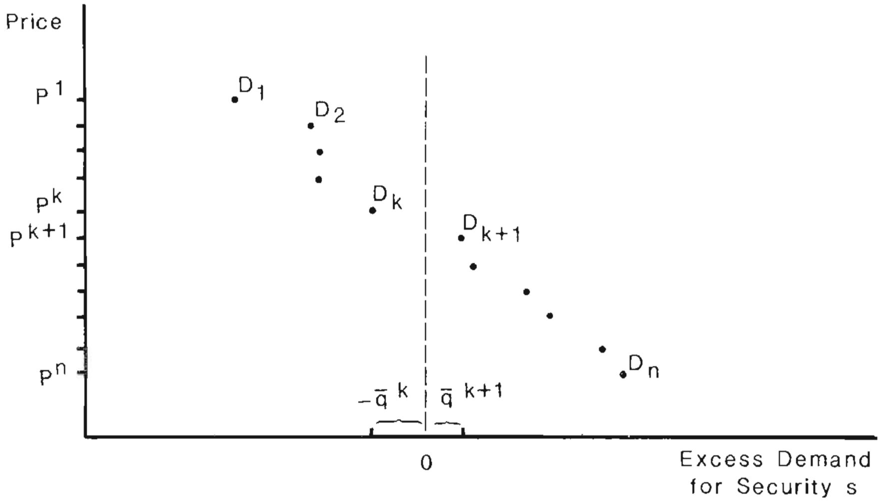

Figure 1. Excess Demand Schedule

at the selected price. The specialist's decision is programmed to follow a strict rule, which prescribes a well-defined course of action for every conceivable excess demand schedule.

A special proviso should be noted here. In extreme cases, it may happen that *all* excess demands for the prices in the relevant range are positive (i.e., that D_1 is to the "right" of 0 in Figure 1) or that all excess demands for prices in the specified range are negative (i.e., that D_n is to the left of 0 in Figure 1). In the former case, price $P_{st} = 1.2\ P_{s,t-1}$ (which corresponds to D_1) will be announced but no trading takes place.[14] This is because the "true" equilibrium price is then clearly greater than 1.2 $P_{s,t-1}$ and intervention by the specialist in filling the excess demand at 1.2 $P_{s,t-1}$ would on balance be disadvantageous since he or she would be selling stock at a price clearly below the (unknown) "true" equilibrium. Similarly, when all excess demands in the range $[0.8\ P_{s,t-1}, 1.2\ P_{s,t-1}]$ are negative, price $P_{st} = 0.8\ P_{s,t-1}$ is "called" but there is no trading.

This procedure has several real-world analogues: the suspension of trading in the presence of news implying sharp price changes is one, and the practice of limiting permissible price changes in a single day is common in commodity markets and can also be found on the Tel Aviv Stock Exchange, for example.

II. Some Alternative Rules of Automated Market Making

When excess demand schedules are not continuous, there will, as noted, be no price for which excess demand is zero (except perhaps by pure chance). Given discrete demand schedules, trading can therefore occur only by rationing some quantities (those of sellers or those of buyers) or by intervention of an additional

[14] In the actual simulation, prices are given in eighths (i.e., 10, 10⅛, 10¼, 10⅜, etc.) in accordance with extant market practices. The quantities 1.2 $P_{s,t-1}$ and 0.8 $P_{s,t-1}$ are therefore rounded (upward) to the nearest ⅛.

party such as a specialist operating as a designated market maker. The present paper is concerned only with the latter possibility. We are interested in operational rules that can govern the specialist's market-making function under a wide range of contingencies. Our purpose here is both to identify rules which appear reasonable on an a priori basis and to examine their properties via simulation.

We are unaware of any developed theories on the subject of demand smoothing. However, if one takes the view that prices should to the greatest extent possible be determined by *investors*, it follows that the market maker's influence should in some sense be "minimized." But in order for his or her impact to be small, it is clearly necessary for the market maker's inventories to be kept low and participation in any given trade to be limited. Consequently, the majority of the rules examined here attempt to make investor demands the overwhelming determinants of prices and trades, by explicitly minimizing some indicator of market interference by the specialist.

The rules which we study fall into two categories. In the first category, the specialist's freedom to determine the trading price is restricted to a very narrow range, and the rules are accordingly labeled "local." Rules in this category are described first. We later describe other rules, where the specialist can choose the trading price in a considerably larger range.

A. *Local Demand Smoothing*

Consider the excess demand schedule (such as in Figure 1) on either side of zero excess demand, i.e., the points D_k and D_{k+1}. One possibility would be to let the specialist pick either D_k or D_{k+1} according to some further specified rule. In the first case, the specialist would buy $-\bar{q}^k$ shares at price P^k, thereby incurring an expenditure of $-\bar{q}^k P^k$ dollars. In the second case, the specialist would sell $\bar{q}^{k+1}$ shares at P^{k+1}, receiving $\bar{q}^{k+1} P^{k+1}$ in compensation. We shall refer to rules which always restrict the specialist to a choice between points D_k and D_{k+1} as *local* demand smoothing rules. In such rules, the specialist seems as though he or she atttempts to *minimize participation in each trade.* The results of our simulation indicate, however, that these rules do not necessarily minimize some *overall measures* of the specialist's interference in the market.

A.1 *The Infeasibility of Price Smoothing*—How might one choose between D_k and D_{k+1}? In view of the emphasis in extant markets on "price smoothing,"[15] one could perhaps consider as a natural candidate for P_t the point which minimizes the absolute change from the previous price P_{t-1}, that is, the price-quantity pair corresponding to

$$\min\{|P^k - P_{t-1}|, |P^{k+1} - P_{t-1}|\} \tag{4}$$

(with ties decided by minimizing $|q_t^0 P_t|$). The problem with this rule is that although the specialist's expected inventory of shares over time is zero, the

[15] For example, one of the demands placed on the specialist is that he or she "... maintains a continuous market with price continuity" (see, e.g., Leffler and Farwell [23]). However, the affirmative obligation to stabilize prices apparently exists only in the U.S. (see Ho, Schwartz, and Whitcomb [21]).

variance of the inventory of shares grows without bound.[16] Thus, the probability that the specialist will eventually hold practically *all* shares, or be short a similarly absurd number of shares, tends to 1. This suggests that any rule that is independent of the specialist's inventories breaks down, a fact already noted by Garman [17] in a related context. Our simulations confirmed this as well; in a typical case, the specialist using rule (4) was, after 250 periods, short more than 66% of the outstanding shares, after having reached a maximum long position of more than 73% in period 176.

We conclude then that when demands change in response to independent events, price stabilization, despite its institutional acceptance, raises not only theoretical questions but also serious practical difficulties.[17] This is not to say that "interference" with respect to price, without regard to inventories, for the purpose of attaining a "true" equilibrium is undesirable or unworkable in *all* contexts. The success of such policies, however, will in general depend critically on the quality of the information that guides them.

A.2 *Some Feasible Local Demand Smoothing Rules*—In the present context, then, any workable rule of interference used by the specialist must pay careful attention to pre-trading position (c_{t-1}, q^0_{t-1}). Table I summarizes a set of local demand smoothing rules that were examined in a market with a single security being traded.

The first rule, LS1, is extremely simple: it instructs the specialist to interfere via a purchase if a long position is not currently held, and to sell if it is. The next two rules, LS2 and LS3, focus on only one of the two quantities c_t and q^0_t, while LS4 and LS5 sum absolute departures from zero of both cash and the value of the portfolio, giving each equal weight in the process. Rule LS6 minimizes the sum of absolute departures from zero of cash, the value of securities, and the net position $c_t + v_t$. Rules LS5a, LS5b, and LS5c modify LS5 by placing constraints on the specialist's maximum absolute position in the security. Thus, in LS5a, for example, the specialist is precluded from being more than 3% long or short; when selection of either D_k or D_{k+1} fails to satisfy this requirement, no trading will occur but a new price is announced according to rule LS5. Similarly, rule LS5d limits participation in any given trade to 50%.

Several of the preceding (single-security) demand smoothing rules were adapted to the case in which several securities are simultaneously traded (see Table II). LM1 is actually LS3 applied independently to each security. LM2 is essentially LS5, with a separate cash account being maintained for each security. LM3 modifies LM2 by equalizing, via redistribution, the cash position associated with each security as soon as trading is completed. Only LM4 is a fully interactive rule: the specialist selects that combination of points from the S pairs $(D_{sk}, D_{s,k-1})$, for all s, which minimized $|c_t| + |v_t|$.

[16] This is because, under our assumptions, increments to the inventories of the specialist would be independent, with zero drift.

[17] The same problem can occur (and did occur) in our simulation when the new price is set using more extreme price smoothing, i.e., by adjusting P_{t-1} (to get P_t) by only a fraction of the price change indicated in Equation (4).

Table I

Local Demand Smoothing Rules, Single Security

LS1	Select $\begin{cases} D_k & \text{if } q^0_{t-1} \le 0 \\ D_{k+1} & \text{if } q^0_{t-1} > 0 \end{cases}$, each t
LS2	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \lvert c_t \rvert$ subject to (3) and (5), each t
LS3	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \lvert q^0_t \rvert$ subject to (1) and (5), each t
LS4	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert q^0_t P_{t-1} \rvert\}$ subject to (1), (3), and (5), each t
LS5	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), and (5), each t
LS5a	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), (5), and (7), $q^0_t \le 0.03Q$, each t. If no solution exists, set $(c_t, q^0_t) = (c_{t-1}, q^0_{t-1})$ and choose P_t as if constraint (7) were absent.
LS5b	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), (5), and (8), $q^0_t \le 0.02Q$, each t. If no solution exists, set $(c_t, q^0_t) = (c_{t-1}, q^0_{t-1})$ and choose P_t as if constraint (8) were absent.
LS5c	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), (5), and (9), $q^0_t \le 0.01Q$, each t. If no solution exists, set $(c_t, q^0_t) = (c_{t-1}, q_{t-1})$ and choose P_t as if constraint (9) were absent.
LS5d	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), (5), and (10), $\lvert \bar{q}^0_t \rvert \le 0.5 \sum_{i\ge 1} \{\bar{q}^i_t \mid \bar{q}^i_t > 0\}$, each t. If no solution exists, set $(c_t, q^0_t) = (c_{t-1}, q^0_{t-1})$ and choose P_t as if constraint (10) were absent.
LS6	$\underset{\{D_k, D_{k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert + \lvert c_t + v_t \rvert\}$ subject to (2), (3), and (5), each t.

Note: $Q \equiv \sum_{i\ge 1} q^i_{st}$. (5) is the boundary price constraint; see [19, fn. 13] for details.

Table II

Local Demand Smoothing Rules, Multiple Securities

LM1	$\underset{\{D_{sk}, D_{s,k+1}\}}{\text{Min}} \; \lvert q^0_{st} \rvert$ subject to (1) and (5), each s and t
LM2	$\underset{\{D_{sk}, D_{s,k+1}\}}{\text{Min}} \; \{\lvert c_{st} \rvert + \lvert q^0_{st} P_{st} \rvert\}$ subject to (2), (5), and (11), $c_{st} = c_{s,t-1} - \bar{q}^0_{st} P_{st}$, each s and t
LM3	$\underset{\{D_{sk}, D_{s,k+1}\}}{\text{Min}} \; \{\lvert c_{st} \rvert + \lvert q^0_{st} P_{st} \rvert\}$ subject to (2), (5), and (12), $c_{st} = \sum_s c_{s,t-1}/S - \bar{q}^0_{st} P_{st}$, each s and t
LM4	$\underset{\{D_{1k}, D_{1,k+1}, \cdots, D_{Sk}, D_{S,k+1}\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), and (5), each t

Note: (5) is the boundary price constraint; see [19, fn. 13] for details.

Table III

Total Demand Smoothing Rules

	Single Security Rules
TS1	$\underset{\{D_1,\cdots,D_n\}}{\text{Min}} \; \lvert q_t^0 \rvert$ subject to (1) and (5), each t
TS2	$\underset{\{D_1,\cdots,D_n\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert q_t^0 P_{t-1} \rvert\}$ subject to (1), (3), and (5), each t
TS3	$\underset{\{D_1,\cdots,D_n\}}{\text{Min}} \; \{\lvert c_t \rvert + \lvert v_t \rvert\}$ subject to (2), (3), and (5), each t
	Multiple Security Rule
TM1	$\underset{\{D_{s1},\cdots,D_{sn}\}}{\text{Min}} \; \lvert q_{st}^0 \rvert$ subject to (1) and (5), each s and t

Note: (5) is the boundary price constraint; see [19, fn. 13] for details.

B. Total Demand Smoothing

Rules which permit the specialist to select any one of the points D_1 through D_n in Figure 1 will be referred to as *total* demand smoothing rules. Such rules clearly give the specialist more flexibility (never less flexibility) than local demand smoothing rules by generally permitting a larger reduction in (absolute) inventories) for any given beginning inventory level. Table III describes a number of total demand smoothing rules. In the single security case, rule TS1 in Table III is a generalization of LS3, while rules TS2 and TS3 are extensions of LS4 and LS5. Similarly, rule TM1 is a direct extension of rule LM1.

III. The Performance of the Simulated Market Maker

Details of the results of our simulation experiments are presented in Tables IV–VIII. In this section, we review the more pertinent aspects of these results, and comment on their implications with respect to the potential for programmed market making in organized exchanges.

A. Local Demand Smoothing

Beginning with the simplest context of unconstrained local demand smoothing for a single security, we found that rules LS1 through LS6 are virtually indistinguishable in most dimensions.[18] On average, the specialist accounted (except under LS1) for 7.7% to 8.0% of all trading; in some trades, however, the specialist was the sole seller or sole buyer. This suggests that his or her role in keeping the market going was rather significant in the thin markets of our experiments. Note also that trading occurred in roughly 39 out of 40 periods (97.5% of the time), that is, prices changed less than 20% in 39 periods out of 40.

[18] The results for rules LS2 through LS4 are omitted from Table IV since they closely resemble those for LS5. LS5 and LS6 *are* different rules, but *they* were similar enough to make them indistinguishable for the series of random numbers generated in the experiment reported in Table IV.

Despite a substantial degree of participation in *trading*, the specialist's average *absolute holding* in the security was 1.28% or less, and did not exceed 5% except under rules LS1 and LS5 (LS6). The specialist's absolute *net* position was even smaller, both on average and at the extremes.[19] Perhaps the most striking thing here is that the differences between LS2 through LS6 are so minor as to be negligible. It apparently makes little difference whether one minimizes the magnitude of the specialist's cash position, share holdings, or combinations of the two.

When limits were placed on the specialist's absolute holdings, all but three of the measures, including those measuring net position, were essentially unchanged (see Table IV). Not surprisingly, average absolute holdings decreased from 0.91% for the unconstrained case to 0.43% when the absolute constraint was 1%. The corresponding absolute maximum holdings decreased from 5.63% to 1%. The specialist's average participation in executed trades also declined from 7.9% to 3.5%. Offsetting these declines, however, was an *increase* in the number of periods in which no trading occurred, from 2.4% in the unconstrained case to 5.2% with the 3% constraint, 8.8% with the 2% constraint, and 24.4% with the 1% constraint. In sum, then, the results suggest that the tradeoffs generated by placing limits on the specialist's absolute position in a security are such that this type of constraint merits serious consideration.

Placing a 50% limit on the specialist's participation on any trade increased the percentage of periods in which no trading occurred from 2.4% to 3.6% (only). However, the absolute security holdings of the specialist increased substantially, from 0.91% to 2.45%, on average, while the maximum increased from 5.6% to 8.1%, as shown in Table IV. The specialist's net position also ended highly unfavorably (−7.1% of the total worth of the security). These results suggest that limiting the specialist's degree of participation in each trade has rather unfavorable consequences on inventory and thus lacks promise as an instrumental variable.

When one specialist handles more than one security, there appear to be three major items to be noted. First, most statistical measures were virtually the same across the four rules that were tested; in fact, the results were essentially unchanged from the single security case (Table V). Second, the absolute net holdings of securities were only about half of those shown for a single security, both on average and at the maximum. This is, of course, not unexpected since long positions in some securities will be offset by short positions in other securities. Finally, the lowest values in both of these categories occurred for rule LM4, the only one to fully consider all securities jointly (this rule also led to a slightly higher degree of average intervention by the specialist, 10.5%). Somewhat surprisingly, the partially joint rule LM3 did worst in terms of the specialist's holdings of securities.

[19] One exception was the case of rule LS1, under which the specialist (inexplicably) made a great deal of money (from a starting position of 0, the specialist ended up with 11.49% of aggregate wealth). We are inclined to view the financial consequences of rule LS1 as a statistical outlier and to view LS1 as inferior to LS2–LS6 on the basis of its greater degree of intervention (it accounted for nearly 11% of the average trading and caused the specialist to hold more stock as well in absolute terms).

Table IV

Comparison of Constrained Local Demand Smoothing Rules

			Bounds on Stock Inventory			50% Bound on Participation
			3%	2%	1%	
Specialist's Rule	LS1	LS5, LS6	LS5a	LS5b	LS5c	LS5d
No. of investors	100	100	100	100	100	100
No. of periods	250	250	250	250	250	250
No. of securities traded	1	1	1	1	1	1
No. of shares outstanding	2,500	2,500	2,500	2,500	2,500	2,500
Initial price	10	10	10	10	10	10
Minimum price	6⅜	6¼	6¼	6¼	6¼	6¼
Maximum price	118½	118¾	118½	118¾	116⅜	118⅜
Final price	32⅝	32⅛	32⅜	32⅜	32½	30⅝
Average price change	4.07	4.00	4.01	4.02	3.87	4.01
% of periods with trading	97.2	97.6	94.8	91.2	75.6	96.4
% of shares traded:						
Average	11.20	11.01	10.79	10.51	9.62	10.96
Maximum	53.56	50.53	72.40	72.13	74.58	51.67
% participation in trades by specialist:						
Average	10.86	7.88	6.26	5.82	3.51	7.33
Maximum	100.00	100.00	46.63	52.59	42.10	42.10
Final position of specialist:						
Cash	9,165	−666	331	−711	−487	−176
Stock	210	−505	−415	−385	−397	−5,254
Net	9,375	−1,171	−84	−1,096	−884	−5,430
Absolute % of stock held by specialist:						
Average	1.28	0.91	0.79	0.71	0.43	2.45
Maximum	7.78	5.63	2.72	1.98	1.00	8.11
Final	0.26	0.63	0.51	0.48	0.49	6.86
Net absolute position of specialist as % of market:						
Average	3.77	0.52	0.85	0.95	0.56	3.67
Maximum	17.91	1.70	2.05	1.78	1.81	7.26
Final	11.49	1.46	0.10	1.35	1.09	7.09

Next we turn our attention to the relationship between the specialist's role in trade execution and the depth of the market. Comparing the first column of Table VI to the first two columns of Table IV, we observe that the specialist's role was considerably greater when there were only 20 investors than it was with 100 investors. As the number of investors was increased from 20 to 50, the specialist's average percent of absolute net holdings decreased sharply, as did net absolute position as a fraction of the total market and involvement in a given trade, while other measures (e.g., average price changes) were essentially unchanged. The specialist's average absolute position also declined sharply as the number of stocks handled increased from 10 to 50 (see middle columns). Again, the joint rule LM4 (see the fifth column in Table VI) did better in terms of the

Table V

Comparison of Multiple Security Local Demand Smoothing Rules

Specialist's Rule	LM1	LM2	LM3	LM4
No. of investors	100	100	100	100
No. of periods	250	250	250	250
No. of securities traded	5	5	5	5
No. of shares outstanding	12,500	12,500	12,500	12,500
Initial prices	10	10	10	10
Final prices	26⅞–35¾	27⅛–35⅞	26¼–36¼	26¾–35⅞
Average price change	4.35	4.35	4.18	4.35
% of periods with trading	96.4–98.4	96.8–98.4	96.4–98.0	96.0–98.0
% of shares traded:				
Average	11.18	11.19	11.18	11.27
Maximum	74.27	86.50	75.97	83.19
% participation in trades by specialist:				
Average	8.45	8.85	8.72	10.46
Maximum	93.13	93.13	86.45	95.38
Final position of specialist:				
Cash	19,764	3,502	19,073	6,661
Stocks	−14	4,286	11,538	2,538
Net	19,750	7,788	30,611	9,199
Absolute net % of stocks held by specialist:				
Average	0.42	0.49	0.78	0.30
Maximum	1.64	2.42	2.53	1.57
Final	0.00	1.11	1.32	0.66
Net absolute position of specialist as % of market:				
Average	1.77	0.54	1.90	0.64
Maximum	9.90	2.60	8.57	2.62
Final	5.14	2.02	3.49	2.40

specialist's holdings than the decentralized rule LM1 (see the second column) where there were 10 securities.

B. Total Demand Smoothing

A comparison of the results for total demand smoothing rules with those of the local demand smoothing rules shows a high similarity not only in fundamentals (price patterns, average price changes, the proportion of periods with trading, and the amount of trading) but also with respect to the specialist's participation in trading and net position (Table VII). As expected, on average the specialist's participation in trading was consistently higher under total demand smoothing than under local demand smoothing for the three rules tested. Furthermore, there was a noticeable decrease in the specialist's absolute percentage holdings of stock for the rules which minimize the magnitude of the stock inventory (LS3, TS1). Otherwise, the differences between rules TS1, TS2, and TS3 appear to be small.

As was the case under local demand smoothing, an increase in the number of investors or the number of securities that the specialist handles had a favorable effect on the performance of total demand smoothing rules. All the measures of specialist intervention dropped dramatically as the number of investors was

Table VI

Comparative Analysis of Local Demand Smoothing Rules

Specialist's Rule	LS3	LM1	LM1	LM1	LM4
No. of investors	20	100	100	100	100
No. of periods	250	250	250	180	250
No. of securities traded	1	10	20	50	10
No. of shares outstanding	500	25,000	50,000	125,000	25,000
Initial prices	10	10	10	10	
Final price(s)	28¼	17¾–41⅞	21⅝–38⅛	14⅛–32⅛	17⅝–41⅞
Average price change	5.62	4.30	4.37	3.98	4.30
% of periods with trading	96.40	95.60–98.40	95.60–98.00		94.80–99.20
% of shares traded:					
Average	14.05	11.14	11.03	10.12	11.29
Maximum	95.54	68.60	73.85	71.19	69.16
% participation in trades by specialist:					
Average	24.47	8.37	8.03	8.74	11.39
Maximum	100.00	97.48	100.00	100.00	100.00
Final position of specialist in $					
Cash	7,350	−3,000	32,308	38,524	19,111
Stocks	−115	753	1,136	−3,699	93
Net	7,235	−2,247	33,444	34,825	19,204
Absolute net % of stocks held by specialist:					
Average	4.21	0.34	0.20	0.13	0.23
Maximum	24.00	1.46	1.40	0.85	1.15
Final	0.82	0.10	0.07	0.14	0.01
Net absolute position of specialist as % of market:					
Average	10.58	0.75	0.65	0.38	0.68
Maximum	52.14	3.52	3.25	1.99	2.73
Final	51.45	0.30	2.12	1.31	2.55

increased from 20 to 100. In addition, as the number of stocks increased from 1 to 50, the average values of both the "absolute net % of stock value held by the specialist," and the "net absolute position of the specialist as a % of the market" declined steadily: from 0.56% to 0.10% for the former and from 0.88% to 0.14% for the latter.

The superiority of the total demand smoothing over the local demand smoothing rules was further documented in a sequence of extensive experiments devoted totally to a comparison between these two classes of rules: the specialist's inventory as a fraction of the market and the specialist's net position as a fraction of the market were both lower. On the other hand, the specialist's participation in individual trades was somewhat greater under total demand smoothing.[20] The

[20] As shown in Table VIII, the specialist's average participation in trades was 12.06% in the case of total demand smoothing and 10.58% under local demand smoothing. It is interesting, although probably strictly coincidental, that the average annual participation rates of specialists on the New York Stock Exchange during the 1977–1981 period ranged from 11.2% to 12.4% (see Wheeler [29, p. 12]).

Table VII

Comparative Analysis of Total Demand Smoothing

Specialist's Rule	TS1	TS1	TM1	TM1	TM1
No. of investors	20	100	100	100	100
No. of periods	250	250	250	250	180
No. of securities traded	1	1	5	20	50
No. of shares outstanding	500	2,500	12,500	50,000	125,000
Initial prices	10	10	10	10	10
Final prices	28	32⅜	27½–36⅛	21⅞–38⅛	14⅛–31⅝
Average price change	5.35	4.02	4.35	4.36	3.97
% of periods with trading	96.80	97.60	96.40–98.00	95.60–98.00	
% of shares traded:					
Average	14.36	11.05	11.28	11.18	10.18
Maximum	95.37	50.51	75.97	73.74	71.18
% participation in trades by specialist:					
Average	27.81	9.29	9.97	9.35	10.04
Maximum	100.00	73.45	86.62	95.03	100.00
Final position of specialist:					
Cash	1,243	−461	2,138	4,222	17,502
Stocks	29	71	−1,685	3,528	−931
Net	1,272	−390	453	7,750	16,571
Absolute net % of stock value held by specialist:					
Average	2.09	0.56	0.35	0.16	0.10
Maximum	20.46	4.74	2.37	1.08	0.59
Final	0.21	0.09	0.44	0.22	0.03
Net absolute position of specialist as % of market:					
Average	2.16	0.88	0.65	0.29	0.14
Maximum	9.28	5.54	1.82	1.35	0.91
Final	9.08	0.48	0.12	0.49	0.62

data yielded at least one statistically significant result (Table VIII): the estimated mean difference between the specialist's absolute inventory under local demand smoothing and inventory under total demand smoothing is 0.3075%. The standard deviation of this variable is 0.076%, or about one fourth of the estimated difference.

IV. The Costs of Demand Smoothing

The essence of the excess demand function depicted in Figure 1 is its discreteness. The role of the specialist (as viewed in this paper) is to intervene by either purchasing or selling a sufficient number of shares for the market to "clear." But this demand smoothing function is not without its cost, quite apart from the purely operational expenses of maintaining the mechanism in question. Consider for a moment the local demand smoothing rules we have examined, each of which involves a choice of either point D_k or D_{k+1} on the part of the specialist. Interpolating the excess demand curve between D_k and D_{k+1}, we see in Figure 2 that it crosses zero at price P^d, where $P^k > P^d > P^{k+1}$. But when the specialist sells $\bar{q}^{k+1}$ shares, he or she does so at $P^{k+1} < P^d$, not P^d; as a buyer, $-\bar{q}^k$ shares

Table VIII

Comparison of Total and Local Demand Smoothing (TS1 vs. LS3)

Average Price Change Per Period	Absolute % of Stock Held by Specialist				% Participation in Trades by Specialist				Net Absolute Position of Specialist as a % of Market				Ending Cash Position	
	TS		LS		TS		LS		TS		LS			
	Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.	TS	LS
4.02	0.56	4.74	0.84	4.38	9.29	73.45	8.01	100.00	0.88	5.54	0.33	1.09	−461	−965
1.85	0.48	4.89	0.77	6.39	16.67	96.97	13.79	95.35	0.38	1.57	0.66	3.09	879	651
3.15	0.48	5.27	0.92	6.07	11.10	85.70	10.59	94.07	1.16	5.96	2.57	15.14	−931	584
0.86	0.45	4.92	0.82	4.55	24.57	96.38	20.65	100.00	2.43	9.25	2.87	10.87	−663	−494
4.67	0.56	4.03	0.68	3.47	8.80	77.69	7.18	97.74	2.18	8.06	1.20	6.94	−3,721	−3,754
2.51	0.53	11.28	0.80	5.53	13.42	100.00	12.48	100.00	4.52	26.16	5.31	30.93	−2,698	−3,074
6.77	0.60	3.84	0.82	3.45	6.70	43.43	5.84	41.57	0.85	7.98	2.96	27.86	−2,557	1,898
8.46	0.66	4.22	0.91	4.60	7.01	80.75	5.98	83.84	3.41	14.10	2.83	9.91	−8,922	5,021
1.81	0.55	7.58	0.87	11.00	16.22	98.93	14.34	99.15	3.13	13.92	4.31	19.70	2,218	2,638
3.49	0.64	5.02	1.00	6.73	12.41	87.49	10.72	100.00	1.38	7.36	3.38	16.83	1,857	−6,712
3.26	0.59	5.95	0.90	6.76	11.40	79.58	10.76	68.83	2.50	25.66	7.84	65.37	929	4,641
6.38	0.53	4.87	0.99	5.38	6.97	61.60	6.00	64.36	0.39	1.73	0.66	2.59	−3,966	−1,767
3.27	0.46	3.48	0.77	4.71	10.45	89.84	8.76	63.21	2.40	11.17	1.43	7.68	−2,342	−933
3.27	0.55	4.07	0.91	4.23	10.02	59.27	8.67	62.17	0.96	5.28	0.62	2.99	1,791	−1,946
1.47	0.59	6.33	0.90	5.55	18.53	100.00	16.26	96.89	0.50	1.83	1.47	6.73	167	−847
3.41	0.47	2.84	0.78	4.48	9.41	68.71	9.21	66.28	1.02	7.11	0.94	7.67	−1,499	2,314
Avg. 3.67	0.55	5.20	0.86	5.46	12.06	81.24	10.58	84.59	1.76	9.54	2.46	14.72	−1,245	−172
Max. 8.46	0.66	11.28	1.00	11.00	24.57	100.00	20.65	100.00	4.52	26.16	7.84	65.37	—	—

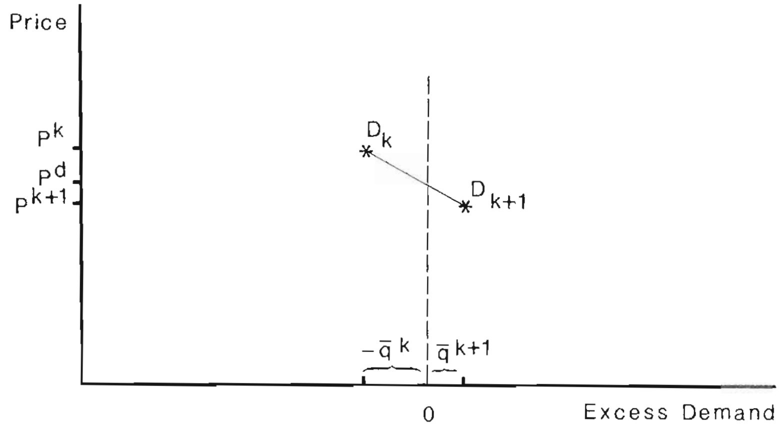

Figure 2. Local Demand Smoothing

are acquired at P^k, not $P^d < P^k$. Thus, assuming that $P^k - P^{k+1} = 1/8$ (since the probability that $P^k - P^{k+1} \geq 1/4$ is very small when the number of investors is large in our model), and considering the specialist's tendency to limit interference by picking the smaller of $\bar{q}^k$ and $\bar{q}^{k+1}$, we see that the specialist would "lose" $1/16$ or less per share bought or sold. In the majority of the simulations, the specialist accounted for about 11% of all trading. Thus, the maximum cost to the demand smoothing function performed by the specialist per share traded to his or her account is approximately $0.11/(0.89 + 1) \times 1/16 = 0.0036$, or $1/3$ of a penny per share traded by investors! The near-negligibility of this cost is confirmed by our simulations: in 56 simulations, the specialist ended with a net gain 26 times and with a net loss 30 times, or what we would have predicted if the smoothing cost had been near zero. Furthermore, as the number of investors increases, the specialist's participation in trading decreases, which in turn reduces overall demand smoothing costs even further.

V. Concluding Remarks

Roughly, our findings may be summarized as follows:

1. Our simulations confirm that, in the presence of independent stochastic increments to (excess) demand, any rule used by the specialist that effectively ignores inventory position, such as any rule dedicated to "price continuity," will ultimately break down.
2. Rules which continually minimize the specialist's absolute share holding or absolute cash position, or weighted combinations of these, yielded surprisingly similar effects on inventory patterns. These effects were overwhelmingly favorable in the sense that the specialist's positions in cash and securities remained "small." In this context, rules which only consider the sign of the specialist's (quantity of) share holdings, or which limit the specialist's activity in any given trade, performed less favorably. However, constraining his or her (absolute) position in securities, while reducing the

number of times trading occurs, appears to present alternatives that merit serious consideration.

3. The average percent of total trading in which the specialist is involved decreased sharply as the number of investors increased.
4. The average net absolute holdings of securities and the average absolute net worth of the specialist as a fraction of the total market, as well as the variability of these measures, decreased sharply with the number of investors and the number of securities handled by the specialist.
5. Total demand smoothing rules appear to be superior to the local demand smoothing rules in terms of their effect on the specialist's absolute inventories.
6. The cost of the demand smoothing function due to buying at a "high" price or selling at a "low" price (but excluding the operational costs) appears to be negligible.

Overall, a number of the programmed demand smoothing rules we examined performed well. In particular, the extremely simple rule that calls for the (computerized) specialist to minimize new absolute share holdings in each security at each trading point via total demand smoothing appears to have noteworthy merits. This rule has the additional property that it is security-wise fully decentralized. In conclusion, we are inclined to view the results of our simulation as showing rather promising potential for automated market making in organized exchanges.

REFERENCES

1. Yakov Amihud and Haim Mendelson. "Dealership Market: Market-Making With Inventory." *Journal of Financial Economics* 9 (March 1980), 31–53.
2. Amir Barnea. "Performance Evaluation of New York Stock Exchange Specialists." *Journal of Financial and Quantitative Analysis* 9 (September 1974), 511–35.
3. Avraham Beja and Barry Goldman. "On the Dynamic Behavior of Prices in Disequilibrium." *Journal of Finance* 35 (May 1980), 235–48.
4. Avraham Beja and Nils Hakansson. "Dynamic Market Processes and the Rewards to Up-to-Date Information." *Journal of Finance* 32 (May 1977), 291–304.
5. ———. "From Orders to Trades: Some Alternative Market Mechanisms." In Ernest Bloch and Robert Schwartz (eds.), *Impending Changes for Securities Markets: What Role for the Exchanges?* Greenwich, Connecticut: JAI Press, 1979.
6. Fischer Black. "Toward a Fully Automated Stock Exchange." *Financial Analysts Journal*, Parts I and II, 27 (July–August and November–December 1971).
7. Ernest Bloch and Robert Schwartz, eds. *Impending Changes for Securities Markets: What Role for the Exchanges?* Greenwich, Connecticut: JAI Press, 1979.
8. Tim Carrington. "Computers, Luck Help Big Board Handle Large Increase in Volume." *The Wall Street Journal* 107 (October 14, 1982), 27.
9. Kalman Cohen, Steven Maier, Robert Schwartz, and David Whitcomb. "The Impact of Designated Market Makers on Security Prices: II, Policy Proposals." *Journal of Banking and Finance* (September 1977).
10. ———. "The Returns Generation Process, Returns Variance, and the Effect of Thinness on Securities Markets." *Journal of Finance* 33 (March 1978), 149–67.
11. ———. "Limit Orders, Market Structure, and the Returns Generation Process." *Journal of Finance* 33 (June 1978), 723–36.
12. ———. "Transaction Costs, Order Placement Strategy, and Existence of the Bid-Ask Spread." Graduate School of Business Administration, Working Paper, Duke University and New York University, 1978.

13. ———. "Assessing the Efficiency of Institutional Arrangements for a National Market System." In Ernest Bloch and Robert Schwartz (eds.), *Impending Changes for Securities Markets: What Role for the Exchanges*? Greenwich, Connecticut: JAI Press, 1979.
14. ———. "Stabilization and Speculation in a Simulated Auction Exchange." Fuqua School of Business, Research Report No. 82-7, Duke University, March 1982.
15. Kenneth Garbade and Chandra Sekaran. "Opening Prices on the New York Stock Exchange." Graduate School of Business, Working Paper, New York University, January 1981.
16. Kenneth Garbade and William Silber. "Technology, Communication and Performance of Financial Markets 1840–1975." *Journal of Finance* 33 (June 1978), 819–32.
17. Mark Garman. "Market Microstructure." *Journal of Financial Economics* 5 (June 1976), 257–75.
18. Barry Goldman and Avraham Beja. "Market Prices vs. Equilibrium Prices: Returns Variance, Serial Correlation, and the Role of the Specialist." *Journal of Finance* 34 (June 1979), 595–608.
19. Nils Hakansson, Avraham Beja, and Jivendra Kale. "On the Feasibility of Automated Market Making by a Programmed Specialist." Finance Working Paper No. 106, Institute of Business and Economic Research, University of California, Berkeley, January 1984.
20. J. M. Hammersley and D. C. Handscomb. *Monte Carlo Methods*. New York: John Wiley & Sons, 1964.
21. Thomas Ho, Robert Schwartz, and David Whitcomb. "Executive Summary of a Comparative Analysis of Alternative Security Market Arrangements." New York Stock Exchange, November 1981.
22. ———. "The Trading Decision and Market Clearing under Transaction Price Uncertainty." *Journal of Finance* 40 (March 1985), 21–42.
23. George Leffler and Loring Farwell. *The Stock Market*. New York: Ronald Press, 1963.
24. Haim Mendelson. "Market Behavior in a Clearing House." *Econometrica* 50 (November 1982), 1505–24.
25. ———. "Analysis of an Exchange Mechanism." Working Paper, Graduate School of Management, University of Rochester, April 1983.
26. Seymour Smidt. "Which Road to an Efficient Market: Free Competition Versus Regulated Monopoly." *Financial Analysts Journal* 27 (September–October 1971).
27. Hans Stoll. "The Supply of Dealer Services in Securities Markets." *Journal of Finance* 33 (September 1978), 1133–51.
28. Seha Tinic and Richard West. "Competition and the Pricing of Dealer Service in the Over-the-Counter Stock Market." *Journal of Financial and Quantitative Analysis* (June 1972).
29. Barbara Wheeler, ed. *Fact Book 1982*. New York Stock Exchange, 1982.