LABORATORY STUDIES IN ASSET TRADING:

PART II--SUPPLY AND DEMAND IN STOCHASTIC MARKETS*

by

Mark B. Garman

August 1973

*This work was supported by NSF, grant GS-32138.

INTRODUCTION

In this work, we deal with simple trading markets in which a single homogeneous commodity is exchanged. Such markets have received considerable attention over the last three-quarters of a century and now form a well-accepted part of almost every economics textbook. Classically, the expositive development of the theory of markets proceeds from the axiomatic basis of consumer (producer) utility functions to the derivation of supply and demand curves and the price-quantity equilibria these imply. In the course of such derivations, certain critical assumptions are made regarding the nature of commodity divisibility, market entry, participant knowledge, and the number of market participants, among others. It is certainly true that imperfections or deviations from these theoretical assumptions, exist in real-world markets. Yet it can reasonably be argued that such imperfections are "removable" in the sense that if they exist in real-world markets, the classical theory is nonetheless valid as a limiting-case approximation. That is, if entry barriers were lowered, if participant knowledge is increased, etc., the imperfections will disappear. This is an important arguement, and we shall take exception to it in only one respect: that the classical theory of markets does not admit of adequate limiting-case analysis in its treatment of time. The nub of the difficulty is that there exists a basic conflict between the joint assumptions of (1) temporal continuity of exchange activity, and (2) its structural stationarity: that as we consider shorter intervals of time (to escape non-stationarity problems), problems of exchange discreteness arise; and as we consider longer intervals of time (to escape discreteness), non-stationarity problems appear. As we shall see later, this conflict is resolved in the classical approach via the implicit assumption that there exist appropriate time intervals big game!

of consideration which simultaneously satisfy both opposing criteria. But we shall assert herein that the difficulty goes deeper than the notion of a (removable) imperfection and necessitates a complete restatement and reinterpretation of portions of economic theory in order to effectively describe exchange activity as it occurs in the real world.

To clarify our general problem statement given above, we shall now examine some of the temporal "imperfections" that seem to exist in actual markets. To do this in a manner which might overcome some of our traditional perceptual biases regarding market imperfections, we shall observe these through the eyes of one contrived creature from the planet Tempus Fugits.

The Tempians, as they are called, are distinguished from human beings by being so small as to be nearly invisible to the latter, and, more importantly, they are possessed of a lifetime of only about one minute Earth time. It has happened that one Tempian recently read an Earthly economics book and has decided to visit a market on Earth. Making the trip in a few long microseconds, he arrives at the New York Stock Exchange, which he understands to be one of the largest centralized markets on that planet. What he finds there is rather surprising to him: He is unable to identify anything that he could term "price" or "quantity" in the sense that buyers and sellers agree upon these in trading. This, of course, is because actual trading in actual markets occurs at discrete points in time: a transaction takes place almost instantaneously relative to the duration of time intervals between transactions. Our Martian friend would have to be lucky indeed to observe one of these rare events in a space of, say, a quarter of his lifetime.

Upon lengthening his visit somewhat, he discovers that other occurrences besides transactions, namely one-sided commitments towards future possible transactions, arise in the market. His textbook has told him nothing about human practices in negotiation and the striking of bargains: the "offer," its acceptance or withdrawal, the completed contract. (Perhaps he will even begin to suspect that these offers themselves, even those not consummated in transactions, may affect market behavior.) Yet he is pleased to find a list, known to us as the Specialists' book, which he interprets as containing supply and demand curves. True, the curves are oddly truncated when graphed:

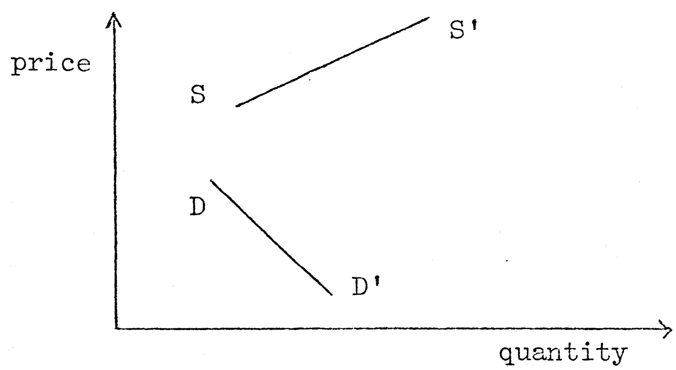

Figure 1

But being faced with a paucity of things he can identify from his textbook, he accepts them as such.

When he finally observes two transactions in a row, he is surprised to discover different "prices" occurring in each, when he is convinced that the slow-moving humans could not have so drastically changed their information sets in such a short (for them) period of time. He is forced to the conclusion that either the parameters of supply and demand are violently non-stationary or else that temporal discreteness in the offers and transactions is itself somehow responsible for the observed phenomenon.

INDIVIDUAL DEMAND

Having dramatized the role of temporal imperfections in actual markets via our Tempian character, we turn now to the nature of the implicit assumptions which are invoked to remove these imperfections in theoretical markets. For convenience and without loss of generality, we shall treat only the demand side of a market.

In the classical approach, individual demand is stated in terms of the quantity of a commodity that a consumer desires to purchase per unit time at a given price. The choice of the time unit length here is of critical importance. If it is chosen too long, the consumer's tastes or budget may change; if it is chosen too short, the demand rate is not a meaningful figure relative to actual (i.e., temporally discrete) purchasing activity in the marketplace. (At what rate per second do you offer by buy applesauce, even assuming you could continuously consume this infinitely divisible commodity?)

But even if the time unit length is chosen to be of an "ideal" intermediate length, a possibility which is implicitly assumed in classical theory, difficulties still remain. For example, consider a consumer who enjoys eating an occasional dill pickle. (To avoid superfluous complications, we shall now assume his desires and his market activity exactly coincide, that is, he desires, offers to buy, buys, and eats a pickle, all in the same instant.) Suppose that the ideal time unit in his case is one month. This means that his tastes may change between but not within months, and that we may make some meaningful statements regarding his market activities within the span of a month. In particular, suppose he consumes 10 pickles per month and derives utility thereby. We argue that it is not only the rate of consumption per month which would yield him utility, but also its pattern. He is likely, for instance, to have strong preferences between a pattern of consuming one pickle every three days during the month and that of consuming ten pickles all on the fifth day of each month, even though the consumption rate per month is the same in each case.

In this paper, we shall attempt to overcome the above difficulties by assuming that individual demand is a discrete stochastic process in continuous time. That is, let $\{N_j(t),\ t \varepsilon [0, \infty)\}$ be a stochastic process, $N_j(t) \varepsilon \{0, 1, 2, \ldots\}$, representing the total number of discrete epochs at which a commodity has been demanded, (i.e., makes a purchase request within some market context) by the jth consumer up to time t. Let $Y_{jn}(\tau_{jn})$ be the amount demanded by the jth consumer at the nth demand epoch, given that it occurs at time τ_{jn}, $\tau_{j,i+1} > \tau_{ji}$, $i = 1, 2, \ldots, N_j(t) - 1$. (Note again that discreteness in time, i.e., $N_j(t) \varepsilon \{0, 1, 2, \ldots\}$, has nothing to with the economic assumption of infinite divisibility of the commodity, i.e., $Y_{jn}(\tau_{jn}) \varepsilon [0,\infty)$). Then

$$X_j(t) = \sum_{i=1}^{N_j(t)} Y_{jn}(\tau_{jn}) \qquad (1)$$

represents the total amount demanded by the jth consumer in the interval $[0, t]$.

The distributions of $X_j(t)$, $Y_{jn}(\tau_{jn})$, and $N_j(t)$ will certainly be conditioned by a number of factors. First, these random variables are assumed to be associated with demands actually presented within a marketplace, and not necessarily with either the fulfillment of those demands or the actual consumption of the commodity by the consumer. We therefore deal in his "wants" or "desires" only insofar as these are made known by the consumer in a context where possibilities for exchange exist, and we specifically admit that actual purchase might not ensue and that consumption is not necessarily synonymous with purchase. Thus $X_j(t)$, $Y_{jn}(\tau_{jn})$, and $N_j(t)$ may be conditioned by variables specific to the consumer: his purchase success, his inventory, his price expectations, etc. Secondly, market considerations will perhaps condition $X_j(t)$, $Y_{jn}(\tau_{jn})$, and $N_j(t)$: market structure insofar as its capacities for "remembering" his demands, his access to information regarding the demand amounts and prices of others, etc. Thirdly, these distributions may be conditioned by prices, since the defined j random variables describe a time series of unilateral "contracts to buy."

There are several arguments for adopting the stochastic process formulation. To begin with, it may not be far from an accurate description of reality. To illustrate this, the reader should inquire of himself as to exactly when he will offer to purchase his next automobile or tube of toothpaste. Even though we like to think of ourselves as "rational" economic agents (either in the loose sense of precise planning, etc., or in the more exact sense of a system of axioms), it would still be difficult to temporally pinpoint such a market activity.

Another argument is that the stochastic process formulation aids greatly the analysis of large ensembles of market participants. To consider an analogue, the physicist would be severely handicapped if he could employ only deterministic models in describing the behavior of a large ensemble of particles: the mass, position, and velocity of each and all interactions of these amongst all particles would have to be taken into account. Indeed, statistical mechanics provides a vastly more simple approach to gaining the aggregate-level insights he pursues.

It is therefore not our goal in the present paper to precisely describe individual consumer behavior, although we could certainly at this point formulate a theory consistent with the stochastic demand process assumed above. (Such a theory would probably include a stochastic income stream, probabilistic budget constraints, axioms for rational choice amongst stochastic consumption streams, and so forth.) But we are concerned rather with aggregate market behavior and shall adopt the attitude of the physicist who cares not whether the individual particles he works with possess rationality, free will, blind ignorance, or whatever, as long as his statistical mechanics will accurately describe the behavior of large ensembles of these particles.

Before passing to aggregate market behavior, however, we should note some of the strengths and weaknesses of the stochastic formulation given. On the plus side, the classical formulation of individual demand represents a special case wherein utility functions are presumed to depend only on the mean value functions of the stochastic processes in question. For example, the mean value function

$$m_j(t) = E[X(t)] \tag{2}$$

could be interpreted as the classical "amount demanded per unit time" if this is stationary over the ideal time interval. (The implicit assumption is therefore that consumers are indifferent amongst all stochastic demand processes having identical mean value functions; compare the pickle example.) On the minus side, the stochastic formulation ignores (or at least complicates the description of) feedback and "gaming" activities of a consumer within a market. In a later paper we hope to explore a game-theoretic formulation (involving strategies again effected at discrete points in continuous time) of individual demand behavior wherein strategies are affected by market conditions.

AGGREGATE DEMAND

We turn now to the consideration of a collection of M individuals, each of which is described by the individual stochastic demand processes defined in the foregoing section. Let

$$N(t) = \sum_{j=1}^{M} N_j(t) \tag{3}$$

and

$$X(t) = \sum_{j=1}^{M} \sum_{i=1}^{N_j(t)} Y_{jn}(\tau_{jn}) \ . \tag{4}$$

Then under some rather weak (for our purpose) assumptions described by Loeve (, p.), the superposition of the otherwise arbitrary counting processes $N_j(t)$ in (3) approaches a Poisson distribution as M increases. Intuitively, these conditions amount to independence in the individual demand processes, disallowing of an infinite number of demands in a finite time period, disallowing the dominance of any one demand process, etc., but their formal description is rather mathematical. A simple analogue is the arrival process of telephone calls to a switchboard: if telephone users act independently according to their own stochastic usage processes, total call arrivals are commonly assumed (and empirically observed) to be well represented by a Poisson process.

The Poisson limit is a testable implication of the foregoing assumptions. It is also a starting point for theories of market structure, which we shall treat in detail in the next report of the current series. Before passing to that topic, we should note a few implications regarding gross market behavior. First, all of the foregoing arguments about individual and aggregate demand processes apply equally well to the supply side, or alternatively, to an excess-demand formulation. Secondly, the mean value function of X(t) is just the sum of the individual mean value functions.

$$m(t) = \sum_{j=1}^{M} m_j(t) \tag{5}$$

Regardless of whether these are stationary or not, we could then depict the familiar price-quantity diagram at time t as:

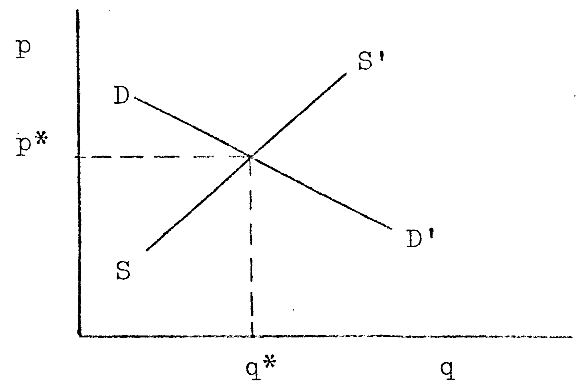

Figure 2

involving two intersecting mean-value functions of price, one each for the aggregate stochastic demand and supply processes, D-D' and S-S' respectively. (We could give analogous arguments for the slopes of the mean-value functions, but this is unnecessary for our purposes.) The point is that now the interpretations are quite different from the classical formulation: The supply and demand curves now give the <u>instantaneous rates</u> of supply and demand rather than absolute supply and demand for some idealized time interval. Furthermore, these rates are "on the average" amounts and do not necessarily have specific physical realizations in the marketplace. We recognize the possibility that there will be sampling fluctuations in actual supply and demand. This leads in turn to a different interpretation of the "equilibrium" point (p*, q*) as it now represents a kind of stochastic equilibrium. Actual prices and quantities under such a form of equilibrium may fluctuate randomly, even under conditions of stationarity.

SELECTION FUNCTIONS

In the above discussion, we have presented rather loosely the stochastic analogues to classical supply and demand curves. For our ultimate purpose of describing market behavior, we shall have to be more precise as to the meaning of such curves. Let us suppose that $\lambda_D'(p,t)$ is the mean value function for some aggregate demand process, i.e., an "instantaneous mean rate of demand" λ_D' which is expressed as a function of price p and time t (cf. formula (5)). Like its classical analogue, it represents a latent demand or "what-if" situation: given that price is forced to an arbitrary value p, the corresponding demand rate should be $\lambda_D'(p,t)$. In classical market theory, it is presumed that p can be and is forced to some (usually the equilibrium) value, hence only one point on a demand or supply curve is "active" at any one time; even though the remainders of the supply and demand curves are presumed to exist, they are depicted as having no effect upon market activities. In the stochastic formulation, this assumption is too restrictive. That is, we shall permit the actual and observable activities of market participants to be affected by perhaps wide ranges of their latent demand and supply curves. (For example, a group of individuals bidding in an auction can offer to buy at any of several prices--with perhaps some restrictions on their range of values, e.g., prices must be in round dollars, each bid must exceed the last by $10 or more, etc. In doing so, they translate their latent desires in market activity.)

We shall deal with this notion mathematically via a so-called aggregate selection function. This will be defined as a function $w(p) \geq 0$ such that the product $\lambda_D(p,t) = w(p)\lambda_D'(p,t)$ represents the actual (observable) market demand activity. That is, $\lambda_D(p,t)$ is the mean value function of the market demand process, and $w(p)$ is an intermediary which transforms latent demand into marketplace demand. The reason for introducing $w(p)$ is so that we may concisely depict the end effects of price knowledge, expectations, etc., upon the way in which a group of individuals presents their demands in a market context. (For example, we don't hear people proclaiming that they would purchase twenty Cadillacs at five dollars apiece, even though they would perhaps do so if they could; their knowledge about Cadillac prices is such that they would not "select" this bid to present in the marketplace.)

The selection function most commonly encountered is the one with a point mass at one price, as in Figure 3:

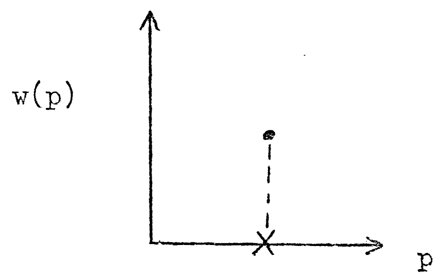

Figure 3

It represents the situation of price-takers, for example price bids for the goods in a department store: here all prices are marked, and bidding less than the marked price almost never occurs since department stores are not known to be flexible in this regard, while bidding more than the marked price is senseless. On the other hand, any of several prices might be bid for shares of some security, as Figure 4 illustrates:

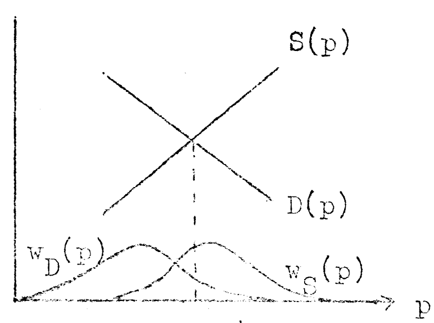

Figure 4

In this example, $w_S(p)$ is the selection function for supply, $w_D(p)$ that for demand; on the average, suppliers are establishing higher prices than demanders, and apparently a continuum of prices is permissable.

In concluding this section, we should note a few observations regarding our selection functions. First, they are dependent not only upon the psychologies of market participants, but also upon the market structure, i.e., the rules of bidding and offering within the marketplace. Secondly, selection functions may have a number of other independent variables besides p, e.g., historical prices, information access, information delay, economic indicators, and whatever else affects individuals' presentation of their "supply and demand" characteristics to the marketplace.

The principal goals of the current report have now been accomplished. We have provided for the existence of discrete transactions, equilibrium price fluctuations, and related phenomena. The next report deals with these phenomena within specific market structures.