1. Introduction

1.1 Overview

Market microstructure is the study of the trading mechanisms used for financial securities. There is no "microstructure manifesto," and historical antecedents to the field can probably be found going back to the beginning of written language, but at some point, the field acquired a distinct identity. As good a starting point as any is the coinage of the term "market microstructure" in the paper of the same title by Garman (1976):

> [W]e depart from the usual approaches of the theory of exchange by (1) making the assumption of asynchronous, temporally discrete market activities on the part of market agents and (2) adopting a viewpoint which treats the temporal microstructure, i.e., moment-to-moment aggregate exchange behavior, as an important descriptive aspect of such markets.

Microstructure analyses typically touch on one or more of the following aspects of trade.

1.1.1 Sources of Value and Reasons for Trade

We generally assume that the security value comprises private and common components. Private values are idiosyncratic to the agent and are usually known by the agent when the trading strategy is decided. Common values are the same for everyone in the market and are often known or realized only after trade has occurred. In security markets, the common value component reflects the cash flows from the security, as summarized in the present value of the flows or the security's resale value. Private value components arise from differences in investment horizon, risk-exposure, endowments, tax situations, etc. Generally, common value effects dominate private value effects. A necessary condition for gains from trade within a set of agents is contingent on some sort of differentiation. In modeling, this is often introduced as heterogeneous private values.

1.1.2 Mechanisms in Economic Settings

Microstructure analyses are usually very specific about the mechanism or protocol, used to accomplish trade. One common and important mechanism is the continuous limit order market. The full range, though, includes search, bargaining, auctions, dealer markets, and a variety of derivative markets. These mechanisms may operate in parallel: many markets are hybrids.

1.1.3 Multiple Characterizations of Prices

The market-clearing price, at least at it arises in usual Walrasian tatonnement, rarely appears in microstructure analyses. At a single instant there may be many prices, depending on direction (buying or selling), the speed with which the trade must be accomplished, sometimes on the agent's identity or other attribute, and the agent's relationship to the counterparty (as well as, of course, quantity). Some prices (like bids and offers) may be hypothetical and prospective.

1.2 Liquidity

Security markets are sometimes characterized by their "liquidity." Precise definitions only exist in the contexts of particular models, but the qualities associated with the word are sufficiently widely accepted and understood that the term is useful in practical and academic discourse.

Liquidity impounds the usual economic concept of elasticity. In a liquid market, a small shift in demand or supply does not result in a large price change. Liquidity also refers to the cost of trading, something distinct from the price of the security being bought or sold. Liquid markets have low trading costs. Finally, liquidity has dynamic attributes. In a liquid market, accomplishing a purchase or sale over a short horizon does not cost appreciably more than spreading the trades over a longer interval.

Liquidity is sometimes defined as "depth, breadth, and resiliency." In a deep market if we look a little above the "current market price", there is a large incremental quantity available for sale. Below the current price, there is a large incremental quantity that is sought by one or more buyers. A broad market has many participants, none of whom is presumed to exert significant market power. In a resilient market, the price effects that are associated with the trading process (as opposed to the "fundamental" valuations) are small and die out quickly.

It is sometimes useful to characterize agents as suppliers or demanders of liquidity. Liquidity supply has traditionally been associated with the financial services industry, i.e., the brokers, dealers and other intermediaries that are sometimes called the "sell side" of the market. Liquidity demanders in this view are the "customers," the individual and institutional investors characterized by "trading needs" (and sometimes called "the buy side").

From a narrower perspective, liquidity supply and demand differentiates agents who are available to trade or offer the option to trade, and those who spontaneously decide to trade. Thus, liquidity suppliers are passive and demanders are active. In any particular trade, the active side is the party who seals the deal by accepting the terms offered by the passive side. In other words, the passive side "makes" the market and the active side "takes."

With the rise of markets that are widely, directly and electronically accessible, the role of liquidity demander or supplier (in the sense of the preceding paragraph) is a strategic choice that can be quickly reversed. The alignment of liquidity demand and supply with particular institutions, therefore, is of diminished relevance in many modern markets.

The *liquidity externality* is a network externality. The attributes of liquidity discussed above are generally enhanced, and individual agents can trade at lower cost, when the number of participants increases. This force favors market *consolidation*, the concentration of trading activity in a single mechanism or venue. Differences in market participants (e.g., retail vs. institutional investors), however, and innovations by market designers, militate in favor of market segmentation (usually called, in this context, *fragmentation*).

The number of participants in a security market obviously depends on features of the security, in addition to the trading mechanism. If the aggregate value of the underlying assets is high, if value-relevant information is comprehensive, uniform and credible, or if the security is a component of an important index, there will be high interest in trading the security. Ultimately, of course, these qualities are determined endogenously with the market

mechanism. But it is common, when emphasizing the exogenous aspects of these attributes to describe a *security* as being liquid or illiquid.

The sources and origins of liquidity are generally what this book and the field are about. They defy simplistic generalizations, but I have found one expression to be particularly thought-provoking: “Liquidity is created through a give and take process in which multiple counterparties selectively reveal information in exchange for information ultimately leading to a trade.” The words are taken from the offering materials for the Icor Brokerage (an electronic swaps trading platform). It is a practical sentiment that resonates throughout much of what follows.

1.3 Transparency

Transparency is a market attribute that refers to how much information market participants (and potential participants) possess about the trading process. Electronic markets that communicate in real time the bids and offers of buyers and sellers, and the prices of executed trades, are considered highly transparent. Dealer markets, on the other hand, often have no publicly visible bids or offers, nor any trade reporting, and are therefore usually considered opaque.

1.4 Econometric Issues

Microstructure data are distinctive. Most microstructure series consist of discrete events randomly arranged in continuous time. Within the time series taxonomy, they are formally classified as point processes. Point process characterizations are becomingly increasingly important, but for many purposes it suffices to treat observations as continuous variables realized at regular discrete times.

Microstructure data are often well-ordered. The sequence of observations in the dataset closely corresponds to the sequence in which the economic events actually happened. In contrast, most macroeconomic data are time-aggregated. This gives rise to simultaneity and uncertainty about the directions of causal effects. The fine temporal resolution, sometimes described as ultra high frequency, often supports stronger conclusions about causality (at least in the *post hoc ergo propter hoc sense*).

Microstructure data samples are typically large in the sense that by most economic standards observations are exceedingly plentiful (10,000 would not be considered unusual). One would not ordinarily question the validity of asymptotic statistical approximations in samples of this size. It is worth emphasizing, though, that the usual asymptotic results apply to correctly specified models, and given the complexity of trading processes some degree of misspecification is almost inevitable. Furthermore, despite the number of observations, the data samples are often small in terms of calendar span, on the order of days or at best months.

Microstructure data samples are new (we don’t have long-term historical data for most markets). The samples may also be characterized as old, though, because market institutions are changing so rapidly that even samples a few years previous may be seriously out of date.

1.5 The Questions

Here is a partial list of significant outstanding questions in market microstructure:

- What are optimal trading strategies for typical trading problems?
- Exactly how is information impounded in prices?
- How do we enhance the information aggregation process?
- How do we avoid market failures?
- What sort of trading arrangements maximize efficiency?
- What is the trade-off between “fairness” and efficiency?
- How is market structure related to the valuation of securities?
- What can market/trading data tell us about the informational environment of the firm?
- What can market/trading data tell us about long-term risk?

Although they might have been worded differently, most of these problems have been outstanding as long as the field has been in existence.

1.6 Readings

This book draws on material from economic theory, econometrics and statistics, and descriptions of existing market institutions. Harris (2003) is a broad treatment of economic theory and trading institutions at the advanced MBA-level. O'Hara (1995) is the standard reference for the economic theory of market microstructure. Brunnermeier (2001) surveys information and price formation in securities markets, treating microstructure in a broader economic context. Lyons (2001) discusses the market microstructure of the foreign exchange market, providing a useful alternative to the present treatment, which is based more on equity markets. Survey articles include Hasbrouck (1996), Madhavan (2000), Biais, Glosten and Spatt (2002). Amihud, Mendelson and Pedersen (2005) survey the rapidly growing field that links microstructure and asset pricing. Shephard (2005) is a useful collection of key readings in stochastic volatility. This research increasingly relies on high-frequency data, and therefore more deeply involves microstructure issues.

Some characteristics of security price dynamics are best discussed in context of the larger environment in which the security market operates. Cochrane (2005) is a comprehensive and highly comprehensible synthesis of the economics of asset pricing. Related background readings on financial economics include Ingersoll (1987), Huang and Litzenberger (1998), and Duffie (2001).

The empirical material draws heavily on the econometrics of time series analysis. Hamilton (1994) is the key reference here, and the present discussion often refers the reader to Hamilton for greater detail. For other econometric techniques (and in particular, duration and limited dependent variable models), Greene (2002) is particularly useful. Alexander (2001), Gourieroux and Jasiak (2001) and Tsay (2002) discuss financial econometrics; Dacorogna, Gencay, Muller, Olsen and Pictet (2001) focus on high-frequency data. The econometric coverage in these excellent books partially overlaps with the present text.

It is difficult to cite authoritative sources covering institutional details of the specific markets. Markets that are recently organized or overhauled, particularly those that feature standard mechanisms, are usually well documented. The trading procedures of the Euronext markets are in this respect exemplary (Euronext (2003)). Hybrid markets that have evolved over extended periods of change and adaptation are much less straightforward. The practicalities of current trading on the New York Stock Exchange, for example, would be extremely difficult to deduce from the codified Constitution and Rules (New York Stock Exchange (2005)). Comerton-Forde and Rydge (2004) provide useful summaries of trading procedures in many securities markets and countries.

1.7 Supplements to the Book

My web site (at www.stern.nyu.edu/~jhasbrou) contains a number of links and programs that may help the reader follow, apply, emend or extend the material in the book. Most of the mathematical derivations in the book were generated using *Mathematica*. The *Mathematica* notebooks are available on the site. Using *Mathematica* does not by any means guarantee the correctness of a derivation, but it does lessen the likelihood of a simple algebraic mistake. A *Mathematica* notebook documents a calculation in standard form. It facilitates the modification and extension of an argument, visualization, and (when necessary) the transition to numerical implementation. The solutions to most of the exercises are contained in the notebooks. The site has several SAS programs that illustrate the techniques.