High frequency data in financial markets: Issues and applications [1]

Charles A.E. Goodhart [a,2], Maureen O'Hara [b,*,3]

[a] *London School of Economics, London, UK*
[b] *Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853-4201, USA*

Abstract

The development of high frequency data bases allows for empirical investigations of a wide range of issues in the financial markets. In this paper, we set out some of the many important issues connected with the use, analysis, and application of high-frequency data sets. These include the effects of market structure on the availability and interpretation of the data, methodological issues such as the treatment of time, the effects of intra-day seasonals, and the effects of time-varying volatility, and the information content of various market data. We also address using high frequency data to determine the linkages between markets and to determine the applicability of temporal trading rules. The paper concludes with a discussion of the issues for future research. © 1997 Elsevier Science B.V.

[*] Corresponding author. E-mail: ohara@johnson.cornell.edu.

[2] Norman Sosnow Professor of Banking and Finance.
[3] Robert W. Purcell Professor of Finance.

1. Introduction

Financial markets operate, during their opening hours, on a continuous, high frequency basis. Virtually all available data sets on market activity, however, are based on discrete sampling at lower, often much lower frequency. There are, for example, on average some 4,500 new quotes for the Dm/$ spot exchange over Reuters FXFX screen page every working day; yet most studies of this market are based on one extracted price per day, or per week. The advent of high-frequency (HF) data sets ends this disparity. In some markets, second-by-second data is now available, allowing virtually continuous observations of price, volume, trade size, and even depths. In this paper, we set out some of the many important issues connected with the use, analysis, and application of high-frequency data sets.

One reason why data sets traditionally were low frequency and discrete was the cost of collection and analysis. In general, only those actions resulting in a (legal) obligation between individuals, e.g. a deal involving a purchase of shares for cash, were written down, and even then the resulting audit trails would normally be retained only a short time. The advent of electronic technology has brought a dramatic fall in the cost of gathering data, however, as well as decreased the cost of the simultaneous transmission of 'news' to physically dispersed viewers. This has changed the structure of markets. The London Stock Exchange, for example, has been superseded by an exchange in which traders observe (common) information over electronic screens, but still trade on a person to person basis over the telephone. There is also a growing tendency for such personal trading to be supplemented by electronic trading, as reflected by the rapid growth in automated exchanges that has occurred in the last decade. These structural changes in trading have important implications for both the availability and interpretation of high frequency data, and we discuss these in Section 2.

The availability of continuous-time data sets presents the problem of dealing with a process which is itself time-varying. The forex market during the Tokyo lunch-hour is quite different in many respects from that in normal Asian trading hours; the market around 08.30 EST (when US news is released) is different from that at other hours. How to deal with such differences in not apparent. For example, should we employ differing scaling systems, e.g. scale by equal amounts of business (variously defined) rather than just linearly by time? Does sampling capture the dynamic nature of the market if the underlying stochastic process is not stable across time? We consider these issues in Section 3.

This sets the stage for Section 4 where we review and survey studies on the statistical characteristics of (continuous) financial market processes. Besides searching for nonlinearities, another major current interest in this field is examining the time-varying volatility in such markets, usually in the context of the growing family of ARCH/GARCH. This vast field has fortunately been recently surveyed by Bollerslev et al. (1992), and so we restrict ourselves to surveying only those studies relating to continuous time market data. At the moment, much of the

empirical work remains quite descriptive, looking at inter-relationships between quote and trade frequency, quote and trade price revisions, price volatility, spreads, etc. One area where empirical results and theory have been more closely connected is in the analysis of equity market specialists. There, much of the literature focuses on how market makers learn from trades, and how this in turn affects prices and quotes. Unfortunately, available foreign exchange quote data are indicative rather than firm, and trade data is virtually nonexistant. Moreover, the theory for quote/spread revision has been worked out rather more clearly for *individual* market makers, than for the 'touch', the best available quoted bid and ask, which, in general, will have been posted by separate market makers. This introduces a number of issues into determining the best approach for analyzing these high frequency data sets.

Another issue of importance is whether high frequency data bases will reveal limitations to the efficiency of markets, thereby providing a way of (legally) making an excess return from trading. The belief that financial markets may exhibit complex, nonlinear dynamics suggest that prior tests, e.g. for unit roots, may have failed to discover more complicated temporal dependencies. Recent research has combined a search for nonlinear relationships, and the use of other predictive techniques, notably neural networks, with an examination of the potential profitability of trading rules. What remains to be seen is whether any clear relationship exists between the trading rules actually espoused by technical analysts and the fractal, nonlinear characteristics uncovered in the data. One obvious advantage of HF data sets is that they provide an adequate basis for the testing of chaos, i.e. *deterministic* nonlinear systems. The evidence now appears to show that, while asset market prices exhibit nonlinearities, they are not chaotic.

Section 5 examines the inter-relationships *between* markets, and how movements in prices become transmitted between associated geographical markets or between markets for related assets, such as futures and spot, option and spot, interest rates and the foreign exchange. Arbitrage opportunities are likely to be seized extremely quickly. It is, therefore, only by looking at the highest frequency, continuous time series that one could observe temporal inter-relationships between markets connected by such (arbitrage) inter-relationships. The paper's final section, Section 6, concludes the paper by setting out a number of issues that remain for future research.

2. Data bases and market structure

Our ability to analyze the working of (financial) markets is limited by the availability of relevant data. Market micro-structure studies have depended on access to high frequency data, and on the use of information technology to store and to process the data sets. For example, the continuous series of Reuters

indicative spot quotes for Dm/$ in the newly available HFDF93 data set contains a huge volume of information. [4] Because these data sets record the second-by-second movement of the market, the microstructure, or minute operational details, of the market is very important. Unfortunately, or perhaps fortunately for those engaged in these studies, the structural form of financial markets varies considerably both between markets and over time as markets evolve. So, the extent to which either the empirical findings or the theoretical concepts can be generalised to other financial markets needs to be explored. This issue has taken on additional importance with the growth of automated exchanges. As Domowitz details in his (1993) analysis of execution systems, there are now over 50 automated exchanges around the world. The rules of such exchanges, and the mechanisms by which they affect price setting and behavior, are only just being investigated by researchers. Yet, as we shall note later in this section, it is these electronic exchanges, and their resultant new data sets, that provide the basis for much future research.

The NYSE is the most extensively studied financial market, but it has a number of idiosyncratic features which make it difficult to generalize to other markets. The NYSE is essentially a hybrid market, combining batch [5] and continuous auctions, a dealing floor and an 'upstairs' mechanism for arranging block trades, a limit order book and a designated monopoly specialist. Descriptions of the operations of the NYSE are found in Hasbrouck et al. (1993), and O'Hara (1995). Moreover, the Tokyo Stock Exchange (TSE) is also a hybrid, with such features as saitori (exchange-designated intermediaries), price limits and mandatory trading halts, see Lehmann and Modest (1994). Yet, while each market differs, there are features in common. All *centralised* exchanges keep records of transactions consummated on the exchange, the price, volume and the counterparties involved, and an estimate of the time of the deal. The names of the counterparties are, however, generally regarded as private and potentially commercially sensitive [6].

---

[4] HFDF93 is a data set containing time-stamped quote information on the $/DM, $/yen, and DM/yen exchange rates for October 1992–September 1993. The data set was provided by Olsen and Associates, Zurich, Switzerland.

[5] For the effect of the opening auction on the NYSE see Stoll and Whaley (1990), and Aggarwal and Park (1994).

[6] Cheng and Madhavan (1994) note, on p. 5, that "it is generally not possible to identify from publicly available databases (e.g. the Trades and Quotes (TAQ) or ISSM databases, whether a trade was directly routed to the downstairs market or was upstairs facilitated. However this distinction is possible with the Consolidated Audit Trail Data (CAUD) files maintained by the New York Stock Exchange. In general, these files are not released publicly; three months (November 1990 through January 1991) of the CAUD files for a sample of 144 NYSE stocks form the basis for the TORQ database which has been widely used in a number of studies." Databases equivalent to TAQ and CAUD are collected by most other centralised exchanges, but the complete audit trail data are rarely released.

The development of high frequency data for other centralized markets has generally been of recent vintage. The Berkeley Options Data base, which dates from the early 1980's, provides time-stamped bid-ask quotes and transaction prices, as well as the current stock price, for each option series traded on the Chicago Board of Options Exchange (CBOE). Because there may be multiple options trading on a single equity, this data allows investigations of the behavior of correlated assets, as well as examinations of the differential behavior of puts and calls. Research using such data is discussed in more detail later in the paper. The data available for US futures markets is even more extensive. Several recent papers (see Fishman and Longstaff, 1992; Chang and Locke, 1996; Smith and Whaley, 1995) have used the computerized trade reconstruction records (CTR) maintained by the Commodity Futures Trading Commission (CFTC). These data include the identity of the floor trader executing a trade, the price and number of contracts in each trade, and the principals behind each trade. Using these data, it is possible to determine which of a floor trader's trades are for customers, personal accounts, and other trading. This information has allowed several interesting papers examining the impact of specific trading rules on one of the largest futures markets, the Chicago Mercantile Exchange (CME). Still, there remain many markets for which such detailed trading information is not available. This is the case in many open outcry markets (such as futures markets) and it is also a common feature of some derivative markets, and of the many markets that employ a batch auction mechanism.

For centralized exchanges, data providing bids and asks (and therefore spreads), and the price and volume of any trade, and the time of each entry [7] is generally available with some degree of accuracy. [8] There are additional data that would be useful for studies of market performance, but are less commonly available. These include information on the supporting schedule of limit orders in an order driven market, or on the change in prices required to persuade market makers to fill an order, in a quote driven system, as the size of the order increases. Such data would allow researchers to construct 'ersatz' demand and supply curves and to study the 'liquidity' and 'depth' of the market. As yet, such data is not widely available.

---

[7] In some cases, as with the release of data on large transactions on the London Stock Exchange, the timing of the *announcement* of the transaction may be delayed behind the timing of the deal itself. Such publication lags may be intentionally intended to influence the availability of information. Whether publication lags are inadvertent, or intentional, their possible presence has to be taken into consideration in any study of market reaction to information. See Board and Sutcliffe (1995).

[8] In most centralised exchanges, a transaction has to 'hit' either the bid or the ask, so one can immediately tell whether it was a purchase (buyer initiated) or a sale. In the NYSE, however, many of the deals are executed within the stated quotes, with a large proportion crossing at the mid-point between the bid and ask. This has given rise to a sizeable literature on empirical studies of the NYSE (see, for example Lee and Ready, 1991 and Petersen and Fialkowski, 1994).

In decentralised markets such as foreign exchange and the interbank money market, there is no quasi-automatic mechanism for providing any information on quotes or trades at all. Participants in these markets are usually fully aware of the current quotes available, but non-bank end users of such markets, e.g. non-bank companies, public sector bodies, are typically not so well informed. Several banks make available information on 'indicative' bid/ask quotes, where indicative means that the bank posting such prices is not *committed* to trade at them, but generally will. These indicative quotes have been collected by the electronic 'news' purveyors, e.g. Reuters, Telerate, Knight Ridder, etc., and disseminated over electronic screens, but they have not typically been archived. Reuters has facilitated and subsidized some researchers (see Goodhart, 1989) to transcribe and make publicly available these indicative quotes for limited time periods. A more extensive data set has been developed and made available by Olsen and Associates, the HFDF93 data, which provides researchers with millions of data points. While very promising in terms of the research questions that can be addressed with this data base, it remains, however, very limited in its coverage, and it contains *no data at all on transactions*.

There are some extremely limited and patchy sources of data on actual firm quotes and transactions on the forex market. Goodhart et al. (1994) obtained access to seven *hours* of data from Reuters electronic broking system, D2000-2, on one day in June 1993, which features firm quotes and transactions. Lyons (1995) has data for the time-stamped quotes, deals and position for a single Dm/$ marketmaker at a major New York Bank, and the time-stamped price and quantities for transactions mediated by one of the major New York brokers in the same market covering a whole *week* in August 1992. Goodhart et al. (1994) concluded that the main characteristics (e.g. the main moments, auto-correlation, GARCH) of the bid/ask series in the indicative data set closely matched that in the 'firm' series, but that the characteristics of the spread in the 'firm' D2000-2 series were distinctly different. The spread in the D2000-2 series was on average lower, much more variable over time, much more auto-correlated, and not bunched at conventional round numbers. Again none of the characteristics of the indicative quote series was a good predictor of transactions. One obvious conclusion is that we need more and better data [9] on 'firm' quotes and transactions from decentralised and OTC markets.

Although the fixed interest, money, bill and bond markets vastly exceed the equity markets in turnover, and may well be of greater macro-economic importance, the number of good market micro-studies in these markets is surprisingly

---

[9] The surveys of the forex market, undertaken by Central Banks under the aegis of the BIS once every three years, now probably to be extended to cover the derivatives market, are extremely useful for some purposes, but are *not* in a format that can help much with market micro-structure studies.

small. Schnadt (1994) examines the UK money market and Goodfriend (1983) and Goodfriend and Whelpley (1986) have done work on the US money market [10], but much of the work on money markets is still descriptive, and the bulk of the empirical work in bond markets still relates to term structure analysis. The absence of much market microstructure analysis in (government) bond markets is particularly surprising since centralised markets in interest rate futures, which can provide associated data, have been established.

A new line of promising research has developed in the area of automated exchanges. While traditional trading venues involve personal interactions between traders either on exchanges or on telephones, the advent of technology permit the development of electronic exchanges devoid of such interactions. As Domowitz (1993) notes, this trend can be seen most clearly in the development of derivative exchanges, where "roughly 82 percent of automated futures/options exchanges have come on line since 1988." Moreover, with only two exceptions, all new derivative exchanges established since 1986 are fully automated, and increasingly new stock exchanges are similarly structured. [11]

The algorithms most automated exchanges employ naturally involve data on price, quantity, time, trader identity, order type, and depths. Dissemination of this information to traders and to outside observers (such as researchers), however, is problematic. In many cases, systems do not display the limit order book even to market participants. The Cotation Assitee en Continu (CAC) in France, for example, has three levels of information, with quotes and trader identification information given only to brokers (see Domowitz, 1993). The availability of data to outside participants and researchers is even more limited. For some markets, outside vendors provide the only access to data, and the extent to which such data is retained (and thus potentially usable for time series studies) is unclear.

Of perhaps equal difficulty is knowing how to interpret and evaluate the data. As noted earlier, most extant theoretical models of market behavior employ variants of an individual specialist who operates in a central exchange. How price formation evolves in automated markets is only now being addressed by researchers. The analysis of Glosten (1994) showing the robustness of an electronic exchange to competition with a market maker system represents a major advance in our understanding of alternative systems. Domowitz and Wang (1994) analyze two computerized market designs with respect to pricing and relative efficiency properties. Bollerslev and Domowitz (1992) consider the effects on volatility of alternative trade algorithms in electronic clearing systems (see also Bollerslev et al. (1994) for an analysis of effects on spreads). Biais et al. (1995) analyse the behavior of the Paris limit order bourse.

---

[10] Also see Pulli (1992) for an excellent study on Finland, and Dutkowsky (1993), for the US.

[11] We thank Ian Domowitz for pointing this out to us in private correspondence.

The variety of structural forms for financial services has allowed some comparisons to be made of the services they provide. In US equity markets, it is common for large trades to transact in the 'upstairs market' where block traders essentially pre-arrange trades. Recent research by Keim and Madhavan (1996) and Seppi (1992) on the differential price behavior of these large trades illustrates an interesting and important application of high frequency data to analyze structural issues. Of perhaps even broader interest is the research investigating the behavior of quote-driven versus order-driven markets (see Pagano and Roell, 1990a,b Pagano and Roell, 1991, 1992, 1995, Madhavan, 1992, de Jong et al., 1993) [12]. This research addresses the important questions of who gains and loses from the resulting price processes in various market settings.

Even within the same trading mechanism, however, there can be large differences in the trade outcomes for different securities. In particular, an area of increasing concern is the pricing behavior of infrequently traded stocks. On the London Stock Exchange, for example, spreads for the most active 'alpha' stocks average 1%, while the spreads for 'delta' stocks average 11%. [13] A similar, albeit much smaller difference, can be found on the NYSE. Why the liquidity of a stock should have such profound effects on spreads is an interesting puzzle. [14] Easley et al. (1996) investigate this problem by using the explicit structure of a microstructure model to estimate the risk of informed trading between active and inactive stocks. Their estimates show that infrequently traded stocks face a higher probability of information-based trades, and hence they argue that the higher bid-ask spreads are necessary to compensate the market maker for the greater risk of trading these stocks. What is intriguing about these results is that they are based on estimates of the market maker's beliefs based on the trade data he observes. As we discuss in later sections of this paper, the issue of learning from high frequency data is fundamental to understanding market behavior, and how this learning differs between market structures is an important topic for future research.

3. The nature of time

Traditional studies of financial market behavior have relied on price observations drawn at fixed time intervals. This sampling pattern was perhaps dictated by

---

[12] Other comparisons have been studied, e.g. floor trading vs. screen trading (Vila et al., 1994), and computerized versus open outcry trading, by Kofman et al. (1994). Also see Benveniste et al. (1992).

[13] Stocks trading in London are divided into four categories based on volume. The most active are called alpha stocks; the least active are the delta stocks.

[14] Amihud and Mendelson (1987, 1988) found that stocks with large bid–ask spreads had higher returns than stocks with smaller spreads. This raises the intriguing, and as yet unanswered, question of whether liquidity is priced in asset markets.

the general view that, whatever drove security prices and returns, it probably did not vary significantly over short time intervals. Several developments in finance have changed this perception. The rise of market microstructure research, with its focus on the decision-rules followed by price-setting agents, delineated the complex process by which prices evolved through time. Whereas prices arising from a Walrasian auctioneer might reasonably have been viewed as time-invariant, prices derived from the explicit modeling of the trading mechanism most assuredly were not. This imparts an importance to the fine details of the trading process, and with it a need to look more closely at the empirical behavior of the market. The concomitant development of transactions (or real time) data bases for equities, options, and foreign exchange provided high frequency observations for a wide range of market data, and hence the ability to analyze market behavior at this more basic level. Finally, the extensive econometric work developing ARCH, GARCH, and related models, which is described elsewhere in this paper, allowed greater ability to analyze this higher frequency data.

A fundamental property of high frequency data is that observations can occur at varying time intervals. Trades, for example, are not equally spaced throughout the day, resulting in intra-day 'seasonals' in the volume of trade, the volatility of prices, and the behavior of spreads. During some time intervals, no transactions need occur, dictating that even measuring returns is problematic. The sporadic nature of trading makes measuring volatility problematic, and this, in turn, dictates a need to view volatility as a process, rather than as a number. These difficulties arise to some extent when the data is drawn on a daily basis, but they become major issues when the data is of higher frequency.

Researchers have dealt with these problems in a number of ways. Brevity requires selectivity in our discussion, so we will focus on only three general issues. These are the implications of clock time versus transaction time, and how this has been handled in the microstructure literature; the mixture of distributions approach to analyzing trade patterns; and the time-scaling approach taken to improve forecasting of security price behavior.

The market microstructure literature attempts to model explicitly the formation of security prices, and hence it seems a natural starting point to consider how the timing of trades affects market behavior. In much of this research, however, time is irrelevant. In the Kyle (1985) model, for example, trades are aggregated and the market price is determined by the net trading imbalance. When the orders were submitted cannot affect the resulting equilibrium. Similarly, while the simple sequential trade model of Glosten and Milgrom (1985) does not aggregate orders, the timing of trades does not convey any information to market participants because time per se is not correlated with any variable related to the value of the underlying asset. In both of these models, only trades convey information, and so the distinction between clock time and trade time is moot.

Diamond and Verrecchia (1987) argued that short-sale constraints could impart information content to no-trading intervals because these constraints might result

in a no-trade outcome when traders would otherwise be selling. Observing a no-trade interval would thus be 'bad news', and prices (and spreads) might be expected to subsequently worsen. This notion of time as a signal underlies the research of Easley and O'Hara (1992). In this model, information events are not known to have occurred, and so the market maker faces the dual problems of deciding not only what informed traders know, but whether there even are any informed traders. In this framework, the market maker uses trades to infer the type of information, and he uses no-trade intervals to infer the existence of new information. Consequently, trades occurring contiguously have very different information content than trades that are separated in time. This dictates that clock time and trade time are not the same.

There are two important empirical implications of this result. First, while prices in the model are Martingales (a property important for market efficiency, an issue discussed later in this paper), they are not Markovian. This has the unfortunate implication that the sequence of prices matters, and hence requires estimation based on the entire history of prices. Second, because time is endogenous, transaction prices suffer from a severe sampling bias and can be viewed as formed by an optional sampling of the underlying true price process. The sampling time is not independent of the price process since transactions are more likely to occur when there is new information. [15] This results in the variance of the transaction price series being both time-varying and an overestimate of the true variance process.

One noteworthy feature of this behavior is that it is consistent with a GARCH framework. GARCH processes can be motivated as resulting from time dependence in the arrival of information, so this model provides an explanation of how such time dependence can occur. A second implication is that volume matters. Since volume is, loosely, inversely related to the time between trades, where the price process goes will differ depending upon whether volume is high or low. [16] The composition of volume will also be important, with expected (i.e., normal) volume reducing spreads, but unexpected volume increasing them. [17] This positive

---

[15] This problem is less serious in the bid–ask quote series because these can be updated by a single individual (i.e. the market maker), while transaction prices await the actions of both an active and a passive party. This suggests that quotes are a better data source (in the sense of being less biased) than are transactions prices. For some markets, in particular FX, only quotes are available, and so analysis of these data may not be seriously affected by these sampling problems.

[16] The dependence of the price process on volume also dictates that volatility will be volume-affected. This issue has been investigated by a wide range of researchers (see for example Lamoureux and Lastrapes, 1990; Campbell et al., 1993; Gallant et al., 1992).

[17] This also implies that volatility will be affected by expected and unexpected volume in a similar fashion.

role for volume contrasts with its standard role in microstructure models where it is largely irrelevant. [18]

That time dependence can affect the stochastic process of security prices is probably not contentious. What is more debatable is *how much* it affects the price process, and this remains inherently an empirical question. Research by Engle and Russell (1995a,b) employs a duration based approach to answering this question. Those researchers explicitly model the intertemporal correlations of the time interval between events. This autoregressive conditional duration model provides an alternative measure to volatility in that the intensity of price changes captures the variability of the order process. The statistical structure of the model provides a framework for testing how the intensity of the price change relates to exogenous variables. The authors find little evidence of such effects in FX data. Other researchers specifically examining the time between trades are Hausman and Lo (1990) and Han et al. (1994), but as yet there are no definite conclusions on the role played by time. A much more extensive literature has developed looking at the links between prices and volume. This literature draws on work by Clark (1973) and Tauchen and Pitts (1983), and it views security prices in the context of a statistical model linking prices, volume, and information.

The mixture of distributions model (MODM) provides an alternative framework for investigating the variability present in high frequency data, and it views the variability of security prices and volume as arising from differences in information arrival rates. The standard model assumes N traders who have different expectations and risk profiles, and these result in different reservation prices. [19] In equilibrium, market clearing requires that the price be the average of these reservation prices. Information arrival causes traders to adjust their reservation prices, and this, in turn, causes trade, which changes the market price. Tauchen and Pitts (1983) assume that these price changes are normally distributed, and this allows them to show that aggregates of price changes and volume of trade are approximately jointly stochastic independent normals. By fixing the number of traders and allowing information events to vary across days, the daily price change and trading volume is then the sum over the within-day price changes and volumes. The Central Limit Theorem can then be used to show that the daily price change and volume can be described by mixtures of independent normals, where the mixing depends on the rate of information arrival.

Fundamental to the MODM approach is that it is new information arrival that changes the reservation prices of trades, and so induces changes in market prices.

---

[18] For example, in the Kyle (1985) model volume is irrelevant because the single informed trader changes his order to offset any volume difference. In the Glosten and Milgrom (1985) model, beliefs are updated on a trade by trade basis, so the aggregate total of transactions provides no information beyond what is already in prices.

[19] This description of the MODM is largely drawn from Harris (1986) and Richardson and Smith (1994).

How exactly the information affects traders, and the related issue of how its dissemination is reflected in trading, is not addressed in this framework. [20] Instead, as Harris (1986) notes, it is assumed that the resulting post information event prices and volume are draws from distributions that are identically and independently distributed for all events. This reflects an interesting contrast with the market microstructure approach, where the focus is precisely on delineating how information affects trading, with prices viewed as the natural outcome of the resultant learning problem on the part of price-setting agents. Whether the statistical approach of the MODM models is a close approximation of the micro-foundations approach of the micro structure literature remains unclear, but both approaches view prices and volume as linked to underlying information events.

The MODM approach can account for a number of regularities in daily data, including heteroscedasticity, kurtosis and skewness in daily price changes, skewness and autocorrelation in daily volume, and positive correlation between absolute daily price changes and volume. Moreover, Nelson (1990) demonstrates that a discrete-time version of the continuous-time exponential ARCH models can be reduced to a MODM, linking these two modeling approaches. Richardson and Smith (1994) argue, however, that much of the evidence supporting the MODM is anecdotal, and that direct testing of the model is complicated by its dependence on unobserved information events. Their analysis finds only mixed support for the model, but their results do suggest some interesting properties of the underlying information flow. In particular, they note that the information flow tends to exhibit positive skewness and large kurtosis. They also show that, while the data are inconsistent with Poisson distributions of information arrival, the lognormal distribution of information event arrivals is consistent with the data. [21]

While this variability in information arrival may, indeed, account for differences in trading throughout the day, there remains the problem of how to analyze the resulting high frequency data. Because the data exhibit 'seasonals', some researchers have employed dummy variables to account for the intra-day variability. While this may be appropriate for some analysis, it does not address the broader issues of why these patterns exist and when they might be expected to be

---

[20] Returns in financial markets, notably on equities but also in forex and bond markets, are much more volatile during hours in which exchanges are trading than when they are shut, at night for domestic markets (other than the international forex market) and over the weekend for all markets. This had been known and reported by several authors, e.g. Fama (1968), Granger and Morgenstern (1970), Oldfield and Rogalski (1980), and Christie (1981), but the salience of this phenomenon was emphasized by French and Roll (1986). Most research on this topic has been done using evidence on returns when some markets were closed, or open, on a particular day (French and Roll, 1986; Barclay et al., 1990), rather than intra-daily data, so we do not pursue this interesting issue.

[21] Such a lognormal modeling approach has been taken by Foster and Viswanathan (1993) to examine volume and volatility patterns in transactions data.

found. What would be particularly useful for the analysis of high frequency data is a blending of the statistical power of the MODM approach with the economic intuition provided by the structural market microstructure approach. The first steps in this direction have been taken in recent research by Foster and Viswanathan (1993), and by Easley et al. (1993, 1995). These researchers use the structure of market microstructure models to analyze the information structure underlying trade data. An advantage of this approach is that it may prove useful in analyzing the properties of intra-day price and volume behavior, an issue clearly of importance in the analysis of high frequency data.

An alternative direction in the treatment of intra-day patterns in trades is a time-scaling approach. The research of Muller et al. (1990) and Muller and Sgier (1992) on the FX markets explicitly recognizes that the time dimension of global trading introduces patterns into trades, and they argue that these patterns, in turn, may be used to determine the expected versus unexpected nature of trade. Their approach, termed the ϑ-time scale, is based on the assumption that there are three main geographic trading areas for foreign exchange. Each geographical area has a particular time pattern, and the global market activity is obtained by cumulating the local patterns. The ϑ-time scale is then computed as an activity measure that essentially expands daytime periods with a high mean volatility and reduces daytime periods (as well as weekend hours) with a low volatility. This time scale allows market activity to be calibrated in a relative sense, thereby introducing an alternative to the clock time and transaction time approaches noted earlier.

This time-scaling approach can also be applied to the analysis of volatility in FX markets. In particular, Muller et al. (1990) and Guillaume et al. (1994) demonstrate that changes in absolute values of spread midpoints (or essentially, the price volatility) can be described by a scaling relation of the form $|\Delta p| = c \Delta t^{1/E}$, where Δt is a time interval and E is the drift component. They argue that this scaling relation holds for a variety of FX rates, and can also be applied to commodities such as gold and silver. Why such a relation holds is not immediately obvious; in spirit it follows early work by Mandelbrot and Taylor (1967) and Fama (1968) investigating the distributions of stock price differences. [22]

Ghysels et al. (1995) combine a time deformation approach with a stochastic volatility model to examine the behavior of FX markets and include both average

---

[22] "a tentative economic interpretation of this scaling law is that it represents a mix of risk profiles of agents trading at different time horizons. The average volatility on one horizon is indeed the maximum return a trader can expect to make on average at that horizon. Alternatively, the average number of directional change for a particular threshold or return is the maximum number of profitable trades a trader can expect to make on average...[T]his relationship between traders with different risk or time-horizon profiles is very stable over the years, notwithstanding the tripling of the volume on the FX markets...[this] is particularly striking as it results from (the fact) that the distribution of the prices changes is unstable and that the conditions of temporal aggregation do not hold", Guillaume et al. (1994, pp. 21–22).

and conditional measures of market activity. The authors use both tick-by-tick data and data sampled at 20 min intervals, which allows them to compare results obtained with different sampling rules. One intriguing finding is that while the geometric average is an appropriate measure of returns on the 20 min scale, it is an unreliable indicator of mean price changes in the tick-by-tick data.

4. The statistical characteristics of intra-daily financial data

4.1. The interaction of volatility, volume and spreads

Perhaps the best known stylized fact about the intra-daily statistical characteristics of the NYSE is that three main features, the volume of deals, the volatility of equity prices and the spread between the bid and ask quotes, all broadly follow a U shaped pattern (or to be more precise, a reverse J). Thus all three variables are at the highest point at the opening, fall quite rapidly to lower levels during the mid-day, and then rise again towards the close (see, among others, Jain and Joh, 1988; Foster and Viswanthan, 1989; Wood et al., 1985; Lockwood and Linn, 1990; McInish and Wood, 1990a, 1991, 1992; Stoll and Whaley, 1990; Lee et al., 1993). A similar pattern tends to hold in other financial markets in which trading cannot easily take place prior to the formal opening. See, for example, Sheikh and Ronn (1994) or Easley et al. (1993) for a study of the daily and intraday behavior of returns on options on the Chicago Board Options Exchange and McInish and Wood (1990b) for the Toronto Stock Exchange. Seasonalities in foreign exchange markets are different and are reviewed in Section 4.7.

The intriguing feature of this temporal intradaily pattern is that it has not proven easy to explain theoretically, at least using the basic model that splits agents in the market informed, uninformed and market maker, (Kyle, 1985; Glosten and Milgrom, 1985; Admati and Pfleiderer, 1988, 1989). Under this latter model one expects uninformed, liquidity traders, with discretion over the timing of their trades, to congregate in time periods when trading costs were low. Given such congregation, and the greater market depth and liquidity that then ensues, privately informed traders also want to trade in such intervals in order to disguise better their identity and information. Nevertheless more information is revealed in such sessions, and thus asset prices are more volatile. On this basis, one can explain a correlation between *volume and volatility* [23], but not at the same time with spreads. As Foster and Viswanathan (1993, 1990) admit "Both the Admati

---

[23] Early papers by Epps and Epps (1976), Tauchen and Pitts (1983) and Bhattacharya and Constantinides (1989) emphasize the role of heterogeneous expectations in influencing the relationship between volume and volatility.

and Pfleiderer, and Foster and Viswanathan models [24] cannot, in their current form, explain the fact that trading volume is highest when trading costs are high for the intra-day tests. While the interday data supports the Foster and Viswanathan model, the use of discretionary liquidity trading in the Foster and Viswanathan model means that it too would predict low volume with high trading costs in an intraday setting" (Foster and Viswanathan, 1993, p. 209).

Equally, a positive association between volatility and the spread, and inversely with depth (Lee et al., 1993), would normally be expected. Greater volatility is associated with the revelation of more information, and with more uncertain markets. If the measure of asset price volatility incorporates the 'bounce' between deals at the bid or the ask, or if when the mean spread is higher, information flows are also higher, then a higher spread will feed back into greater volatility. This finding of a positive correlation between volatility and spread holds for all the micro-structural empirical studies of which we are aware, with the main direction of causality running from volatility to spread rather than the reverse.

So, the peculiar feature of the NYSE that needs special explanation is why the *volume* of deals is so high at the start and end of the trading period. Indeed, as we show later in this Section, the particular U shaped feature in the NYSE for volume does not generalize over other markets. Thus on the London Stock Exchange, where SEAQ does *not* have a formal opening and closing, the pattern of volatility and spreads remain U shaped, whereas "trading volume has a two-hump-shape rather than a U shape over the day" (Kleidon and Werner, 1994). In so far as intra-day quote frequency provides a reasonable proxy for the intra-day volume of deals on the forex market, then there are no signs at all of a U shape in deal volumes in US trading hours (rather the *reverse*), and only rather limited signs of this in Asian and European trading hours (again largely influenced by the lunch-hour dip in quote frequency (see Demos and Goodhart, 1992).

This concentration of volume, at the formal opening and close, has been best modeled by Brock and Kleidon (1992), who extend the model of Merton (1971) to show that transactions demand at the open and close of trading will be both greater and less elastic than at other times of day. Since information about fundamental stock prices [25] and, hence, optimal portfolio proportions will have been varying continuously during the market's closure, there will be a strong demand to trade. Similarly, when prospective market closure foreshadows an inability to readjust

---

[24] The Foster and Viswanathan model differs from that of Admati and Pfleiderer in its assumptions about the temporal pattern whereby an asymmetric information advantage accrues to some investors over the course of the week and is then dissipated by a general public announcement.

[25] One surprising lacuna is that no one seems to have examined whether the characteristics, e.g. volatility, volume, spread, of the NYSE opening are much influenced by the time-varying form of the public news announcements made shortly before the market's opening. This is symptomatic of the seemingly small influence of *public* news announcements on asset prices and the resulting paucity of academic studies of such relationships.

portfolios for 17 1/2 h overnight and over 60 h on Friday night, it will focus investors' attention on the need to rebalance before the closed period arrives. [26]

The release of such pent-up demand to rebalance portfolios thus generates an increase in both (expected) volumes and volatility, as the (market) orders reveal both private and (interpretations of) public information. In view of such higher volatility, the rise in spreads would be expected, whatever the market micro-structure. In their model, Brock and Kleidon also emphasize the monopoly position of specialist traders on the NYSE, and their ability to maximise monopoly profits by cashing in on the increased, and inelastic, demand for transactions services at the open and close. However, it has not been demonstrated empirically that the increase in spreads on the NYSE is significantly greater than on other asset markets with more competitive structures. So, the origin of the peaks in spreads in the NYSE has not yet been clearly identified.

Be that as it may, the relationship between the volume of trading and volatility is quite complex, and depends in some part on whether the fluctuation in volume is expected (i.e. an intra-daily seasonal) or unexpected and 'news' related. Volatility and spreads can be high when markets are thin, as for example over the week-end or in the Tokyo lunch hour on the foreign exchange market. By contrast, when markets become very active, volatility and spreads are also positively correlated. Kim and Verrechia (1991) model volume as the product of the absolute mean change in price from period 1 to period 2, (a measure of the extent of new information being revealed) and an aggregate measure of the heterogeneity of differences of view about such information. Their model provides insight into the relation of volume and public information.

Blume et al. (1994) provide an alternative model of the price–volume–information linkage. In their model, volume is related to the quality of traders' information. This quality linkage arises because traders receive signals of the asset value, where each signal is drawn from some underlying distribution. The precision (or signal quality) of the distribution may be unknown to some traders, reflecting the difficulty of evaluating the quality of any new information. As is standard in noisy rational expectations models, prices reflect the level of the information signal, but the dual nature of the uncertainty precludes a revealing equilibrium using the information in price alone. But because volume is not normally distributed, it can incorporate information that is not already impounded in the price. In this model, volume itself becomes informative, and traders watching volume can know more than traders who watch only prices. Blume et al. demonstrate that this provides a basis for technical analysis of volume data. For our purposes here, this model

---

[26] This argument suggests that the opening and closing peaks in volume, volatility and spreads, should be somewhat attenuated where stocks are cross-listed, so that rebalancing can take before and/or after the primary market closure. Kleidon and Werner (1994) do not find any attenuation in closing peaks in volatility, volume or spreads in London for UK firms cross-listed in the US, nor much significant difference in their opening pattern in the US. See also Breedon (1993).

predicts how both the dissemination of information, and its precision affect the price–volume relation. Moreover, because volume depends on both the quality and quantity of the information signal, the relation of volume to information is nonlinear, and so too is the relation of volume and price volatility.

Karpoff (1987) has surveyed studies of the relationship between volumes and price changes, and has shown that, in equity markets (though not, for obvious reasons, in forex markets) volume rises more when prices are rising, than when they are falling, and that volume is positively correlated with volatility, as measured by the absolute value of price changes. He concludes that, "It is likely that observations of simultaneous large volumes and large price changes — either positive or negative — can be traced to their common ties to information flows (as in the sequential information arrival model), or their common ties to a directing process that can be interpreted as the flow of information (as in the mixture of distributions hypothesis)."

4.2. The determinants of the spread

Whereas explanation of the intra-daily temporal pattern of relationships between volume, volatility and the spread has proven problematic, analysis of the determination of the spread in isolation has been an example of micro-market structure work, theoretical and empirical, at its best and most successful. The theoretical literature focuses on analysing the factors influencing a single market maker is his determination of the spread. Three main factors are identified. First, inventory carrying costs create incentives for market makers to use prices as a tool to control fluctuations in their inventories. Amihud and Mendelson (1980), Zabel (1981), Ho and Stoll (1983) and O'Hara and Oldfield (1986) formally model the effect of market maker inventory control on prices. Second, the existence of traders with private information, the adverse selection motive, implies that rational market makers adjust their beliefs, and hence prices, in response to the perceived information in the order flow. The literature on this includes Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987, 1992), Glosten (1989) and Admati and Pfleiderer (1988, 1989). Third, there are the other costs and competitive conditions which help to determine the mark-up that the single market maker can charge. These conditions are frequently taken as being constant over the day, but in some models, e.g. Brock and Kleidon (1992), can be time varying.

To estimate the empirical effect of these factors, it is helpful to get data of the quotes charged by a single market maker and an estimate of her inventory. This can be done either by examining a market, such as the NYSE, with a single specialist market maker, provided one can make a rough estimate of the volume of deals intermediated by that market maker, or alternatively by having direct access to the books showing the quotes and inventory positions of individual market makers (Lyons, 1996; Neuberger and Roell, 1992; Madhavan and Smidt, 1991, 1993). One problem such models face is disentangling the inventory and informa-

tion effects, since both predict that prices will move in the same direction as order flow, but for different reasons. This difficulty is further compounded by the fact that information-based models generally assume risk neutrality, while inventory models require risk aversion. One approach to deal with this is to start with a general statistical model and then impose certain theoretical restrictions on the coefficients that allow the underlying structural relationships to be identified. Thus Madhavan and Smidt (1991) combine an inventory model with a model of information adjustment (also see Neuberger (1992) for a similar exercise using data from the London Stock Exchange). Such approaches allow empirical estimation of inventory and information effects, but are subject to the criticism that the theoretical restrictions are too severe; in particular, that they rule out any covariance effects between inventory and information.

Previous studies of inventory control by market makers in equity markets (Hasbrouck, 1988; Madhavan and Smidt, 1991; Hasbrouck and Sofianos, 1993; Snell and Tonks, 1995) have found relatively weak intraday inventory effects, though Lyons (1995) [27] found strong inventory control/effects in his model of the forex market. Madhavan et al. (1994) in their empirical study of equity prices suggest "that inventory effects are manifested towards the end of the day, so that the conclusion of these (previous) studies may be worth reinvestigating."

Although the empirical studies provide a diversity of findings, nevertheless our general impression is that these exercises have been ingeniously devised and broadly successful. But such studies do depend either on a single market maker structure or on having data from individual specialist(s). Otherwise with differing market makers setting bid and ask quotes, the spread is not a choice variable, but is endogenously dependent on the decisions of two, usually unidentifiable, separate market makers whose inventory positions are unknown. In the forex market, the FXFX series do show the identity of each bank inputting individual quotes, but not only is their inventory position not known, but also the indicative nature of such quotes makes their use to proxy the underlying market spread a potentially hazardous exercise (Goodhart et al., 1994). Moreover, these models require a simplicity of structure that may not be realistic.

Bollerslev et al. (1994) employ a different approach in their analysis of spread behavior in the interbank foreign exchange market. They develop a methodology that permits characterization of the stationary conditional probability structure of quotes in a screen based system. Their analysis uses the information available to screen traders to estimate how order flow parameters affect spread behavior. This largely statistical approach does not incorporate asymmetric information issues or

---

[27] Oddly enough, when Lyons tested for time-of-day effects he found that inventory control coefficients became muted rather than amplified at the close. He suggests that "One possible explanation of this is that it is precisely at the end of the trading day that marketmakers least want to signal their position via quotes, preferring to trade away from positions through brokers or other marketmakers' prices."

inventory concerns, but it does allow for the stochastic nature of order arrivals and cancellations to influence the level of spreads.

A second empirical approach in the literature to analyzing the spread is to regress the spread on a variety of explanatory variables. A recent example from the forex market is Bollerslev and Melvin (1994), though the caveat about indicative quotes should be remembered. Their main finding is that spreads rise as volatility increases, that ''Measuring exchange rate volatility as the conditional variance of the ask price estimated by an MA(1)–GARCH(1,1) model, we find that there is a strong positive relationship between volatility and spreads.'' Given the link between spreads and volatility, we turn next to a review of how this variable is modelled in high-frequency, intra-daily studies.

4.3. Volatility and memory

In high frequency studies, as in empirical exercises using lower frequency data, the use of GARCH to model the auto-correlation in market volatility still reigns supreme. As described in the survey by Bollerslev et al. (1992) there has been a, still steadily increasing, number of variants developed to catch, inter alia, possible asymmetric effects of large (rather than small) shocks, or price declines (as contrasted with increases, etc.). Nevertheless GARCH, in one or another of its variant forms, is now used almost routinely to model the time path of volatility in almost all studies of financial markets.

There are, however, alternative ways of modeling time-varying volatility; two approaches in particular should be mentioned. The first is to model variance as an unobserved stochastic process (see for example Jacquier et al., 1994; Harvey and Shephard, 1993; Harvey et al., 1994). The second approach is to use the implicit forecast of volatility derived from the option market to forecast subsequent volatility in the spot market (see Harvey and Whaley, 1992; Canina and Figlewski, 1993; Jorion, 1994; Bank of Japan, 1995). The option forecast has often, but not invariably, compared well with a GARCH estimate as a predictor of future spot volatility. What has yet to be shown is whether there are any *identifiable, systematic* factors driving implied option volatility additional to that modelled by GARCH. Comparisons of implied option volatility (relative to GARCH) have so far used lower frequency daily data. There are doubts whether option markets are sufficiently developed to allow for meaningful variations in intra-daily implied volatility to be derived. This prompts the question whether the use of high-frequency, intra-daily data has yet had much *particular* influence on the study of volatility.

There are, perhaps, two main respects in which it has. First, the relative frequency of observations, as compared with identifiable shocks, is much greater in high-frequency series. Somewhat paradoxically, this means that the higher the frequency of the data, the easier it is to study long memory characteristics, although the power of tests may be a problem. And it is here that problems with

the standard GARCH emerge. It is common to find that the coefficients in the standard GARCH equation sum to approximately one in such empirical exercises, i.e. implying IGARCH behavior. That volatility is thus a random walk, and can drift out to infinity or zero, is not intuitively appealing. Most of us tend to believe that volatility should, in the long run, revert to a mean level dependent on the likelihood of natural shocks. But assuming the coefficients sum to *less* than unity requires that the effect of a shock to volatility declines exponentially; and that has been found to give an excessively quick decay rate in several 'long-memory' studies (Ding et al., 1993; Dacorogna et al., 1993). One potential solution is to use Fractionally Integrated GARCH, or FIGARCH, models, which allow mean reversion but at a much slower hyperbolic rate than in GARCH models (Baillie et al., 1993; Baillie, 1994; Baillie and Bollerslev, 1993, 1994), although this technique has yet to be applied to high frequency data.

The other characteristic that distinguishes intra-daily from lower-frequency studies of volatility is that there is much stronger intra-daily seasonality. Much of this intra-daily seasonality in volatility arises from time-of-day phenomena, e.g. market opening and closing (especially in equity markets), the (differential) effect of lunch hours (in the forex market), and the Pacific gap between the close of US markets and the opening of Austral/Asian markets. Such phenomena are akin to standard seasonal effects at lower frequencies, and similar problems apply. One is that, within a Koyck lag framework such as GARCH, entering seasonal dummies along with the lagged dependent variables implicitly assumes that the rate of decay of a foreseen seasonal shock to volatility is exactly the same as that of an unforeseen shock. It is not clear that this should be so in reality. Indeed Andersen and Bollerslev (1994) and Guillaume et al. (1995) show that, unless the (deterministic) intra-daily affects on volatility are taken into account, GARCH coefficients are likely to be spurious, and, even when they are incorporated, GARCH processes often tend to be unstable and unsatisfactory when used on intra-daily data.

An even more acute problem is caused by announcements of economic data, such as the latest figures for the money supply, trade, or inflation. The exact time of such announcements is often known, and considerable effort is made by economists to predict both such figures, and the market's likely reaction. For the most part, this largely known, but time and day-varying, schedule of known 'news' announcements has been ignored in GARCH studies. Guillaume et al. (1994) document that there is a major spike in forex volatility at 08.30 EST, when many of the main US data series are announced. There is, at least, one study (Goodhart et al., 1993) that suggests that standard GARCH coefficients do *not* remain robust when news occasion variables are entered. Ederington and Lee (1993) examine the effect of scheduled US economic news announcements on US interest rates and on the Dm/$ exchange rate using data from futures markets. They find that "[W]hile most of the price changes occurs within one minute, volatility remains considerably higher than normal for another fifteen minutes or so and slightly higher for several hours. This can be explained as either continued

trading based on the initial information or as price reactions to the details of the release as they become available."

Perhaps the most serious problem of GARCH modeling is that we do not yet have a good theory to explain such persistence. Whereas theory can provide a good explanation of the correlation of volume and volatility, it cannot yet explain persistence without either imposing strong restrictions on the sequential process of trades *or* by assuming an unexplained and undocumented persistence in the arrival of information. Lamoureux and Lastrapes (1990) suggest that (the auto-correlation of) volume of trades may be a good proxy for (the auto-correlation of) the arrival of information. Laux and Ng (1993) argue that these results may be contaminated by simultaneous equation bias and, instead, use data on the number of price changes as their proxy for information arrival. Both studies leave unanswered the questions of what are these information arrivals that jointly cause volume and volatility, and why do they exhibit such persistence?

Without theory, it can be argued that GARCH is just a successful method of data fitting. There are currently attempts to apply learning models to explain persistence (see Brock and Le Baron, 1993), but these are beyond the scope of this survey. Perhaps such persistence may also depend upon the heterogeneity of agents with differing operational time horizons. [28] It is not clear why persistence should continue, in the standard informed/uninformed/market-maker paradigm, beyond the elapse of time necessary for current information to be revealed in price changes; and it is not clear why this should lead to long memories and slow decay rates. Equally, while one can see intuitively that the presence of a variety of particular agents (some of whom may have operational horizons of no longer than a few hours, whereas others may have operational horizons extending to quarters or even years), could lead to much greater persistence, it has yet to be rigorously or convincingly formalized and modelled. What does seem a common factor in empirical studies is that it takes volume to drive price volatility. We, therefore, turn next to an examination of the literature on what orders and activity have most effect on prices.

4.4. Which trades move prices?

Other things being equal, a larger trade should be associated with a larger price effect, both because of its effect on the market maker's inventory position and

---

[28] The belief that such persistence is related to heterogeneity is tested by Hogan and Melvin, who add term(s) related to the standard deviation of survey responses (MMS) on forthcoming US Trade Balance data to a GARCH model of conditional volatility on foreign exchange rates subsequent to that announcement. Their results give "evidence that heterogeneous expectations are a source of meteor shower (persistent) effects in the subsequent foreign exchange market. However, for the next two subsequent markets, we find no evidence. Thus our results suggest that heterogeneous expectations can lead to volatility spill-over effects (meteor showers) but that the persistence of such activity is quite limited" (Hogan and Melvin, 1994, p. 245).

because it may imply more confidence by the purchaser/seller of the accuracy of her information. The better informed the trader, the larger the amount that trader would wish to deal at any given price. Easley and O'Hara (1987) demonstrated that this would result in market prices differing with trade size, with large trades occurring at 'worse' prices [29] (also see Holthausen et al., 1987). But ceteris paribus does not always hold. In particular, an investor with private information would like to trade in such a way as to disguise his identity as privy to private information (Laffont and Maskin, 1990). The attempt of (privately) informed investors to hide among the uninformed forms the basis of the theoretical work of Admati and Pfleiderer (1988, 1989) and Foster and Viswanathan (1990, 1993).

Barclay and Warner (1993) argue that most informed trades will be undertaken through medium sized trades, at least on the NYSE. Small trades would take too much time and have excessive execution costs, [30] whereas really large (block) trades are likely to become visible. Trading a large block 'upstairs' on the NYSE often requires some pre-arrangement, during which time information leaks out (Cheng and Madhavan, 1994). Consequently, there is a preference to restrict such trades to uninformed, 'sunshine' trades, Keim and Madhavan (1996). Very large orders, based on private information, are more likely to be broken up and insinuated, more circumspectly, into the market. Barclay and Warner describe this as 'stealth trading'. There is some evidence that one class of informed traders, corporate insiders, do concentrate their orders in medium sized trades (Jaffe, 1974; Meulbrock, 1992; Cornell and Sirri, 1992). Barclay and Warner test for trade size informativeness by examining the ratio of medium sized trades (to small and medium sized trades) using *daily* data during the run-up to tender offers or other events causing systematic unusual behaviour. In our view such hypotheses could be more efficiently tested using intra-daily data.

Easley et al. (1995) use such intra-day data in their testing for size effects. Their approach uses trade data to estimate directly the information content of different trade sizes. If the market maker believes informed traders are more likely to trade larger than smaller amounts, then the probability of information-based trade will be greater for large trades. Such an outcome would be expected if the market is in a separating equilibrium, whereby the profit to an informed trader is larger trading a big quantity at a worse price as opposed to a smaller quantity at a better price. The market can, however, be in a pooling equilibrium, where the informed essentially spread out across trade sizes. In this equilibrium, both large and small trades can be information-based, and trade size effects can be minimal. The authors empirical results find evidence of significant, but varying, trade size

---

[29] This may be mitigated if transactions costs decline with quantity (Black, 1986; Easley and O'Hara, 1987). Bessembinder (1994) suggests that spreads are a negative function of expected volume, but a positive function of unanticipated volume.

[30] Block et al. (1994) examine execution costs on the NYSE by the trading division of a major bank, Nationsbank, and conclude, p. 174, that "larger trades do not have higher indirect execution costs."

effects. While for some stocks in their sample there is no trade size effect at all, in general large trades have twice the information content of small ones.

One intriguing result in this research is that it is transactions, rather than volume, that moves markets. Such an outcome dictates that the rate of trade, rather than the volume of trade, underlies the adjustment of prices. This role of transactions is developed in more detail by Jones et al. (1994). This empirical study of the determinants of price movements finds that volume adds little explanatory power beyond that conveyed by the transaction per se. The role of trades thus emerges as an important area for future research on the price–volatility–information relationship. Of course, in foreign exchange markets it is not possible to study the question of which trades move markets, unless one can get a supplementary database on trades. Lyons (1994) has such a data set, and he has used it to examine the interactions between quote and dealing intensities and price changes. His main finding is that "trades occurring when transaction intensity is high are significantly less informative than trades when transaction intensity is low." He ascribes this to the 'Hot Potato' hypothesis that most deals involve inventory rebalancing among dealers. It remains to be seen if foreign exchange trades, rather than volume, can provide insight into other empirical regularities.

4.5. Auto-correlations and cross-correlations in returns, quotes and trades

The movement of prices following a trade is of obvious importance for understanding the behavior of markets. In the standard sequential trade framework (see Glosten and Milgrom, 1985, a market maker sets new trading prices equal to the conditional expected value of the asset. The subsequent trading prices form a martingale and, thus on an ex ante basis, prices and thus returns should be uncorrelated. If the market maker cares about inventory, however, price changes may be more complex, and in particular may exhibit negative serial correlation due to the market maker's efforts to move his inventory position in a desired direction. If the data do *not* allow complete differentiation between buy orders at the ask and sell orders at the bid, then the first order negative auto-correlation of returns will be accentuated by the bid-ask 'bounce' (Roll, 1984). Evidence of such negative auto-correlation would be more visible the higher the frequency of the data.

In electronic markets, or in specialist markets permitting limit orders, price movements may be affected by the clearing of orders against existing orders. In particular, a large order may move along the limit order book, and/or transact with a number of competing market makers. Rather than display the negative first order correlation in returns, trades and quotes noted above, this can result in positive auto-correlation in these variables. Again such effects would, one would expect, be more prominent the higher the frequency of the data.

As elsewhere, there is more empirical evidence on auto-correlations in the NYSE than elsewhere. Due to the nature of the data, most studies have been undertaken on a trade by trade basis without knowing precisely whether the active

side of the trade was a buy or a sell. Consequently such results incorporate some bias due to the 'bounce' between the bid and the ask (see for example Porter, 1992 and Harris, 1986). After taking account of this effect, the latest findings suggest quite strong signs of *positive* auto-correlation in *trades* (i.e. a trade at the ask is more likely to be followed by another at the ask) (see Huang and Stoll, 1994; Madhavan et al., 1994; Easley et al., 1995; Hasbrouck, 1991a,b, 1988; Hasbrouck and Ho, 1987, and (relatively much weaker) in *returns*, Hasbrouck and Ho, 1987; Lo and MacKinley, 1988). The auto-correlation of trades varies, however, according to whether the stock has a low trade volume, in which case the negative auto-correlation implied by inventory control effects reappears, or a high volume, in which case positive auto-correlation dominates (Hasbrouck, 1988). Hasbrouck (1988) suggests that this positive auto-correlation arises because the NYSE combines a limit order procedure, where one would expect positive auto-correlation, with a specialist, where one would not [31]; thus this positive autocorrelation "is perhaps a consequence of the relatively greater importance for these stocks of public limit orders and relatively lesser importance of specialist transactions."

In the forex market, the only available time series providing data on trades and quotes are the very short series for small and possibly unrepresentative parts of the market obtained by Lyons (1993a,b), Lyons (1994) and by Goodhart et al. (1994). The latter report very strong positive autocorrelation in trades (buys following buys), but approximate random walk in returns. There is, on the other hand, now a huge data set available of Reuters FXFX indicative *quotes*. At intervals shorter than ten minutes, or on a tick by tick basis, these show strong signs of a first order moving average negative auto-correlation (Goodhart, 1989; Goodhart and Figliuoli, 1991; Goodhart and Giugale, 1993; Baillie and Bollerslev, 1990a,b). Most authors ascribe this to the indicative nature of the FXFX quote series, with quotes shifting backwards and forwards between banks with differing order imbalances, persistent tendencies to quote high or low (Bollerslev and Domowitz, 1993), or differing information sets (Goodhart and Figliuoli, 1992). Goodhart et al. (1994) report, however, that negative auto-correlation in quotes remains present in their short, partial series of firm quote data. Goodhart and Payne (1995, forthcoming) ascribe this, along the lines of the theoretical analysis of Ho and Stoll (1983), to the existence of 'thin' markets, so that when the best quote is removed by a trade, the next best price is some distance behind that.

Whatever the reasons, the empirical findings of *very* high frequency autocorrelations (strong positive in trades; *perhaps* weak positive, after taking account of the bounce, in returns; and negative in quotes) are an interesting feature of high frequency data series. Another interesting inter-relationship is that between trades

---

[31] The free option problem of a limit order (see O'Hara, 1995) might also induce such behavior, since keeping a static quote may increase the odds of being hit by an informed trader. Moving prices up and then back may elicit information on trading sensitivities and reduce this problem.

and quotes (and hence spreads as well), which has been developed in pioneering work by Hasbrouck (1991a,b) using data from the NYSE. He finds that the full impact of a trade on the price occurs with a protracted lag, and that as a function of trade size, the innovation on the quote is non-linear, positive and increasing, but concave. Further, he finds that spread size exhibits a response to trading activity, with large trades associated with a widening of the spread. Moreover, trades occurring when spreads are relatively wider have a greater impact than when spreads are narrow. Intriguingly, he argues that the price impact and (by implication) the extent of the information asymmetry appear more significant for firms with smaller market values.

Goodhart et al. (1994) in a similar study using FX data find that knowing the quantity involved in each trade added little (in their case nothing) to the information obtained from the *direction* of trade, a result consistent with the earlier mentioned works of Easley et al. (1995) and Jones et al. (1994). In contrast to Hasbrouck, Goodhart et al. found *no* significant effect of quote revision on order flow, since the *frequency* of quote revision, rather than the size of each quote revision, appears to be the crucial variable determining the likelihood of future trades.

4.6. Seasonality, nonlinearities, neural nets and chaos

One variable that has received considerable attention in the forex empirical work is the frequency of quote entry. Time-varying frequency and irregular spacing of entries is a major feature of the strong (deterministic) seasonal patterns corresponding to intra-daily, daily, and geographical (e.g. Asian, European and American market space) factors in the foreign exchange markets. As noted in Guillaume et al. (1994), "These seasonal patterns are found for the volatility (Bollerslev and Domowitz, 1993; Dacorogna et al., 1993), the relative spread (Muller and Sgier, 1992), the tick frequency (Demos and Goodhart, 1992; Muller et al., 1990), the volatility ratio and the directional change frequency [32], (Guillaume et al., 1994)." The strength of these intraday effects dictates that failure to adjust for them can result in misleading statistical analysis of high frequency FX data.

Adjustments to do so have involved using seasonal dummies (see Baillie and Bollerslev, 1990a,b), time-scaling (see Dacorogna et al., 1993), or some type of Fourier transform (see Andersen and Bollerslev, 1994). This latter adjustment explicitly deals with the nonlinearities found in the data in that a Fourier series is a series expansion of a nonlinear function, generally, in terms of sines and cosines.

---

[32] Guillaume et al. (1994, p. 4) state that "the directional change frequency measures the average number of price changes of a fixed-amplitude over the data sample, that is the time-subintervals are the varying parameter. Using threshold values as a measure of the risk one is ready to take is quite natural to traders."

A related, but more topical approach, is that of a neural network. Neural networks are essentially a nonparametric technique to approximate a nonlinear relationship. As Lo (1994) explains, in common with a Fourier series, a neural network can be viewed as a series expansion of a nonlinear function. The mechanics of a neural net, which are beyond our focus here, are described in detail by Lo (1994). Such neural networks have been used to predict security price behavior (see Trippi and Turban, 1992), option price dynamics (see Hutchinson et al., 1994), and foreign exchange price movements (see Pictet et al., 1992).

Neural networks have a number of interesting implications for dealing with high frequency data sets. Although neural nets are described as 'learning' from the data, this is too strong; like many statistical approaches, neural nets involve updating from the data and revising. Provided the underlying structure of the network is well defined, the neural net can do an excellent job of predicting the statistical movement of prices. But the underlying structure is crucial; as Lo (1994) points out, without sufficient structure, it is not even possible to identify the processes of interest in some classes of problems. For the financial market applications used thus far, it appears that neural nets do provide a valuable means for dealing with the underlying nonlinearities in the data. But just as standard forecasting techniques rely on past data and relationships, so, too, do neural nets. If the underlying structure of the market (or of the price process) changes, then neural nets will fare little better than standard models in predicting out-of-sample values. This suggests that the better our understanding of the underlying structural factors driving price-setting in markets, the better will be our ability to build and apply neural nets to financial markets. This synergy underscores the value of improving extant theoretical models of financial asset price dynamics.

The availability of intra-day quote data in the FX market over a sufficiently long time period has made this market the preferred habitat for studies of nonlinearities, fractal structures, and chaos (see Brock, 1991). Yet explorations for evidence of low-dimensional chaos in FX have all come up with negative results (Vassilicos, 1990; Vassilicos et al., 1992; Vassilicos and Demos, 1994; Guillaume et al., 1994). Abhyankar et al. (1995) suggest that this reflects the power of using the large data sets available with high frequency data; studies based on smaller data sets have been unable to test the sensitivity to initial conditions, a fundamental issue in chaos.

4.7. Technical analysis and market efficiency

The existence of nonlinearities in a market might be expected to allow for *potentially* predictable regularities. If such regularities were themselves sufficiently complex, controversial, time-varying or unstable, then one can envisage the possibility of abnormal returns being made by some investors *without* a rational efficient market competing them away. Until a few years ago, the efficient markets hypothesis, even in its weakest form, was believed to rule out the possibility of obtaining abnormal returns from analysis of the past asset price movements. More

detailed analysis of the micro-structure of asset markets, however, has given rise to a greater willingness to believe in the possibility of certain complicated regularities, and with it the possibility that some of the trading rules of the technical analysts might be applicable. Thus Brock et al. (1992) use a bootstrap technique to test whether some simple moving average strategies can generate abnormal returns. They conclude, pp. 1757–1758, that, "The recent studies on predictability of equity returns from past returns suggest that the conclusion reached by many earlier studies that found technical analysis to be useless may have been premature." (see also Levich and Thomas, 1993; Le Baron, 1992, 1994[33]; Goodhart and Curcio, 1992[34]; and Pictet et al., 1992).

Given the widespread use of technical analysis in financial markets, especially in foreign exchange markets (Taylor and Allen, 1992; Menkhoff, 1995), it is good that economists are now prepared to assess this methodology seriously, rather than just dismiss it out of hand. But these findings do raise two questions for future research. First, why and how can such predictive power arise and persist? Definitive explanations for this are lacking, but there are some promising directions being pursued. The recently developed area of behavioral finance attempts to show how psychological factors can contribute to market prices under- or over-reacting to new information (see for example, Lee et al., 1991). Another possible explanation lies in the underlying learning problem confronting agents. Blume et al. (1994) demonstrate a role for technical analysis of volume data due to the multi-faceted nature of new information. In this model, market statistics beyond price are informative to traders because there is more to be learned than simply the magnitude of some simplistic signal. If new information is imprecise, or if there is uncertainty about the dispersion of new information among traders, then market statistics may impound information in ways not available from price.

This suggests that research on what is conveyed by market data (other than prices) may yield explanations for a number of seemingly anomalous market regularities. While intuitively appealing, this is not an easy task. Most of our theoretical predictions of market behavior are derived in a simple environment where all random variables, including prices, are normally distributed. This construction provides tractability, but it also means that knowing one variable (the price) essentially tells you all that can be known. Departing from this to consider

---

[33] Le Baron shows that "There is reliable evidence that simple rules used by traders have some predictive value over the future movement of foreign exchange prices", but "much of such profitability occurs in those periods when the Federal Reserve is actively intervening...[T]his predictability puzzle is greatly reduced if not eliminated when days in which the Federal Reserve was actively intervening are eliminated."

[34] Instead of constructing artificial trading rules, Goodhart and Curcio use intra-daily data and measures of support and resistance points, provided by Reuters over FXNB, to test for profits from sales/purchases contingent on the market breaking through such points.

volume, for example, means giving up the normal structure, and then the ability to solve, or even characterize, the resulting model is dramatically limited.

A second question raised by the success of technical analysis is what does it mean for a market to be efficient? The exact meaning of efficiency has been contentious from the outset, but this is even more complex an issue in the high frequency environment we consider here. Over short intervals, it is not apparent that even the notion of public information is well defined. Moreover, even if it were well defined, does semi-strong form efficiency really measure anything of interest? Might it not be of more interest to measure or compare the rates at which different market prices impound new information? These issues are well beyond the scope of our analysis here, but we raise them to illustrate that the ability to analyze high frequency observations raises more than methodological issues; viewing markets as continuous processes may require rethinking many of our underlying precepts as well.

5. Inter-market linkages

Studies of inter-market relationships form the second main bloc of empirical exercises in the area of micro-market studies. These have been mainly carried out for equity and equity derivative markets, owing perhaps to the difficulty, until now, of obtaining intra-daily time series for other markets. Unlike studies of individual equity markets, there has been little theory to guide empirical studies of inter-market relationships. The efficient market hypothesis implies that mispricing, and associated arbitrage opportunities, between related markets should be rapidly eliminated, and that can be tested. There are some (theoretical) conjectures whether trading on the basis of news should appear first in a derivative, or its associated spot, market, i.e. the lead–lag relationship; and there are a variety of arguments, on either side, about whether the existence of derivative, especially futures, markets increases, or reduces, volatility in the associated spot market.

Perhaps the most direct investigations of linkages is found in the research on lead–lag relationship between movements in prices (and volatility) between equity and derivative markets. Black (1975) first suggested that the higher leverage available in option markets might induce informed traders to transact in options rather than in stocks. This let loose a torrent of research investigating whether options led stocks, or vice versa. Here the importance of data frequency for empirical study is vividly illustrated. The early research (see for example Manaster and Rendleman, 1982) used daily data and concluded that options lead stock. Vijh (1988) argued that this relation was spurious due to end of day differences in markets, and research with intra-day data by Stephan and Whaley (1990) concluded that, in fact, stocks lead option. Chan et al. (1993), however, argue that this relation is also flawed due to different price discreteness rules in the two markets. Adjusting for these, they argue that neither market leads the other.

While these analyses investigate the relation of option prices and equity prices,

a different approach is taken by Easley et al. (1994). These authors argue that if traders are trading options on the basis of new information, then particular option volumes might have information content for future price movements. In particular, traders would either write a put or buy a call if they knew good news, and would write a call or buy a put if they knew bad news. Signed volume, therefore, could provide information on the future movement of the equity price. Using intraday option data, they show that these option volumes do, indeed, lead equity prices, a finding they interpret as evidence that option markets are a forum for information-based trading.

There are similarly motivated studies of the lead–lag relationship in *returns* between stock index *futures* and the underlying cash market, and in *volatility* between the same two markets. Chan (1992), building on the prior work of Finnerty and Park (1987), Ng (1987), Kawaller et al. (1987, 1990), Harris (1989), and Stoll and Whaley (1990) is a good recent example of the first. Chan et al. (1991) is a good example of the latter. In general, there is a two-way relationship between the markets, both for returns and volatility. Movements in either market can help to predict the other.

In the case of returns, however, the relationship is asymmetrical, with strong evidence that the futures (both Major Market Index and S&P 500) leads the cash index, and much weaker evidence that the cash index leads the futures. Moreover, Chan (1992) concludes that "the lead–lag pattern varies consistently with the extent of market-wide movement. When there are more stocks moving together (market-wide information), the feedback from the futures market to the cash market is stronger. This provides some support for the hypothesis that the cash and futures markets do not have symmetric access to information."

In the case of volatility, Chan et al. (1991) suggest that the intermarket relationships are more symmetrical, concluding that "for the most part, price changes occur simultaneously in both markets... Our evidence is consistent with the hypothesis that new market information disseminates in both the futures and stock markets and that both markets serve important price discovery roles."

To summarise, there is strong evidence of instantaneous, two-way causation between the spot and the associated derivative markets. For returns, the derivative market may lead and have the stronger impact on the spot market than vice-versa (see Huang and Stoll, 1994, who find, using data at five minute periodicity's, "that past stock index futures returns predict subsequent stock returns.").

A second set of studies investigates linkages between geographically separate but related markets. The experience of the 1987 Crash reverberating through world stock markets, and the continuing experience of co-movement between returns in similar markets, e.g. bond and equity, in differing countries has led to a raft of studies either measuring or attempting to explain such. These studies have concentrated on the international transmission, between similar domestic markets, of returns and volatility. Taking the exercises examining returns first, a key paper is King and Wadhwani (1990). They hypothesize that investors will use a signal

extraction model to separate the globally important 'news' in the price changes in the other national market(s), which were open while their market was shut, from the nationally local 'news'. There are, however, many other studies on this subject around the same time, notably Bennet and Kelleher (1988), Von Furstenberg and Jeon (1989), Becker et al. (1990), Schwert (1990), Hamao et al. (1990, 1991), Lin et al. (1991), Neumark et al. (1991), and Susmel and Engle (1994).

In general these early papers found the following stylized facts, as reported in Lin et al. (1991, p. 2). "(I) Volatility of stock-prices is time-varying. It rose considerably around October 1987, but quickly decreased afterwards. (II) When volatility is high, the price changes in major markets tend to become highly correlated. (III) Correlations in volatility and prices appear to be asymmetric in causality between the United States and other countries. The US movement affects other markets, but not vice versa."

With their focus on the relationship between New York, Tokyo, London, and the main continental stock markets, the basic data needed for these studies could be reduced to the opening and closing values on each Stock Exchange, plus, in some studies, the effect on London of the opening of the NYSE, and vice versa (see for example Susmel and Engle, 1994). It is only now in subsequent studies of markets in the same, or nearly overlapping, time zone that finer intra-daily data becomes necessary. One problem that has remained is that opening prices may reflect stale quotes from the previous close, and/or other peculiarities (Lin et al., 1991, also see Aggarwal and Park, 1994).

The results generally uphold the signal extraction hypothesis, that investors do try to extract globally important 'news' from price movements in similar markets elsewhere, and use this to update prices at or after the opening. There is still controversy about how long it may take for such global news to be incorporated in domestic prices, and whether spillovers are symmetric or asymmetric with New York. Our concern, however, is exactly what is this global 'news' that moves international markets. King et al. (1993, 1990) examine a factor model of asset returns and risk premia and conclude that the main common factors driving world stock markets are 'unobservables', rather than public 'news' such as US unemployment data. Similarly, Lin (1994) shows that the Nikkei overnight returns react strongly to movements in the S&P 500, but hardly at all to specific US public 'news', whether in regressions including, or without, the movements in the S&P 500.

Indeed, one puzzle in the study of asset markets, either nationally or internationally, is that so little of the movements in such markets can be ascribed to identified public 'news'. In domestic (equity) markets this finding is often attributed to *private* information being revealed. But in the international context, how could private information be expected to have a 'global' impact?

The second main area of empirical research in this sub-field has been international correlations between markets in *volatility*. Here the canonical paper is by Engle et al. (1990), examining the volatility in the yen/dollar exchange rate. They

distinguish between volatility clustering effects which are specific to one financial centre (a heat-wave) from those which travel between centres (a meteor-shower). They find "that the empirical evidence is generally against the heat wave hypothesis. This rejection is consistent with market dynamics which exhibit volatility persistence possibly due to private information or heterogeneous beliefs, or with stochastic policy coordination or competition." Ito et al. (1992) examine whether 'meteor showers' might have arisen from stochastic policy co-ordination; they look at volatility spillovers in sub-period regimes each with a varying extent of international coordination. They find that this latter cannot plausibly be held responsible for the meteor-shower effect. This again underlies another of the main themes/puzzles that arises from our survey, notably the question of what factors are responsible for the observed persistence in volatility, here often shown to hold between national asset markets, as well as within them. [35]

6. A research agenda for high-frequency data in finance

The ability to access and analyze high frequency data bases provides enormous potential for furthering our understanding of financial markets. As this survey makes clear, the range of questions that can now be addressed is extensive, and it is certainly within the realm of possibility to argue that our current analyses will soon be viewed as arcane relics of a financial 'dark age'. It can also be argued, however, that these new data raise more questions than they settle, and that without some resolution of these problems the use of high frequency data will provide little of value. The task for researchers and practioners alike is to determine from this wealth of data the underlying fundamentals that drive market and asset price behavior. We have raised a number of issues connected with the use of high frequency data, but there are undoubtedly many more that we have not addressed. Even within this limited framework, however, a number of key issues emerge, and it is to these that we now turn our attention. Our goal is to suggest areas where increased study might produce high returns.

A central difficulty in analyzing high frequency data is our lack of knowledge of the underlying theoretical structure of trading. As we have discussed, the paradigm typically employed in analyses of equity markets emphasizes the role of private information, and it models the price process as resulting from a trading game played by informed traders, uninformed traders, and a market maker. Analyses of foreign exchange markets have focused more on public information, and have generally used econometric, rather than structural, models to motivate empirical testing. In this approach, trader heterogeneity plays an important role,

---

[35] Susmel and Engle (1994) cast doubt on the length of such persistence among equity markets, and Hogan and Melvin (1994) test, with mixed results, whether such persistence was greater when prior expectations of 'news' were more heterogeneous.

but where this comes from and why and how this affects prices is not well specified.

These approaches provide valuable insights into market behavior, but their limitations are both apparent and severe. Models with private information suffer from the difficulty that information arrival must be exogenously assumed, so that the timing of information events is specified. In the Glosten and Milgrom and Kyle models, an information event is assumed to occur, and trading then forces prices to true levels. What is not allowed is the realistic possibility that new information could arrive before this adjustment is complete, or that information events are essentially random. The difficulty is that, if there are informed agents with different information, then both trades and the prices at which they execute are informative. The learning models developed thus far, however, can not be solved with such dual uncertainty. While some authors have avoided this problem by assuming that private information is revealed publicly after every trade (see Madhavan and Smidt, 1993 or Admati and Pfleiderer, 1988), this characterization is not generally attractive.

Analysis of public information is, perhaps, even more problematic. The current informed trader paradigms are not well suited to analyzing this problem, and what analyses have been done have generally involved little formal modeling. There is little formal analysis of how trader heterogeneity affects trading behavior, or, by extension, market variables of interest such as prices, returns, or volatilities. There is no theory of maximizing behavior to suggest why certain regularities should arise, or even to provide structure for empirical work. While econometric analyses have provided intriguing insights, it is hard to understand empirical correlations without some underlying economic explanation.

It might be fruitful to analyze more fully the learning problems confronting agents in a high frequency trading environment. That traders learn from watching the market is apparent to even the most casual observer, but what conveys the most information, or what underlies the 'feel of the market', or even what variables in the market trigger traders' desires to trade is not apparent. Determining the role of market statistics is a first step in understanding this process, but this, in turn, rests on our ability to characterize more fully how information in general affects trading. This, as yet, remains undone.

Perhaps a simpler starting point is to consider the curious empirical finding that public information announcements have little impact on some markets. For markets such as FX, public information is surely fundamental, yet empirical research has failed to show its expected impact. Why this is so is surely an issue for future research. The ability to access higher frequency data may provide answers, and an intriguing glimpse of how this can be useful may be found in the research currently done in Accounting. Accounting research traditionally investigates how public announcements by firms affect equity values. The development of high frequency data has allowed for intra-day 'event studies', whereby the impact of announcements is investigated in the hours surrounding its release.

Greater availability of FX data can facilitate similar studies in foreign exchange markets. This would allow investigation of intriguing questions such as how market reactions differ across trading locales, how news is transmitted across markets, and whether the impact of government or central bank announcements differs across countries or markets.

As we have discussed throughout, empirical analysis of market data has revealed a number of conundrums in the market. Among the most puzzling issues is the behavior of volatility. While the general properties of volatility remain elusive, perhaps the most intriguing feature revealed by empirical work on volatility is its long persistence. Such behavior has sparked a search, almost akin to that for the Holy Grail, for the perfect GARCH model, but the underlying question of why such volatility persistence endures remains unanswered. We conjecture that the ability to analyze higher frequency data may be particularly useful in pursuing this issue, in part because it provides an opportunity to investigate longer (in transaction time) time-series of data.

What may be equally useful, however, is the ability to investigate volatility in related markets, and between markets. In particular, volatility is surely the driving force in derivative markets, and much interest in volatility research derives from its obvious applications in the pricing and trading of derivatives. High frequency data on derivatives trading, combined with similar data on spot market trading, would allow researchers to investigate volatility in several dimensions.

A serious omission in the work surveyed above is the lack of research on microstructure issues in fixed income markets. Despite the size and importance of these markets, there are virtually no empirical or theoretical microstructural analyses. We suspect that much of the reason is the lack of adequate data, as the dispersed nature of the market does not lend itself to centralized data collection. Nonetheless, the same technological factors which have permitted the development of FX data bases will surely also facilitate data collection in these markets. Given the rudimentary state of current work, even a small scale study would be of great interest.

Of all the issues revealed by our survey, however, perhaps none is more important than the simple question of what is gained in our investigations of market behavior from moving to higher and higher frequency data observations. While it is natural to think that with respect to data 'more is always better', this need not be so. High frequency data bases are still expensive to collect, maintain, and manipulate. High frequency observations are also subject to a wide range of idiosyncratic factors such as non-synchronous trading, intra-day seasonal effects, measurement errors due to bid–ask spreads or reporting difficulties, and even conceptual problems, such as defining a 'return' during an interval in which no trading occurs.

Two research questions here seem particularly important to us. First, how do the properties of the data set differ with respect to different sampling rules? As we have noted, issues of time take on added importance when examining the

micro-behavior of markets. If, however, observations of the data drawn every minute, or every five minutes, or even every hour have the same statistical properties, then analyzing second-by-second data may not be necessary for every problem. Similarly, if intra-day seasonals reflect nothing more than noise in the overall movement of the underlying true value process, then sampling at lesser frequencies may well be preferable. Without a better understanding of the fine details of the data process, however, it is impossible to know whether sampling the data makes a difference. We believe resolving this issue would be a major step in the use and understanding of high frequency data sets in financial analysis.

A second, and related, question is what information do we want to include in the development of future high frequency data sets? The advent of technology has transformed such a question from idle speculation to a practical necessity, but its answer requires much thought and study. Certainly, it is tempting to conclude that we should simply include 'all' of the information, allowing researchers to pick and choose amongst the information depending upon their particular research topic. [36] But what exactly is 'all' of the data? As recent research has shown, it is not just prices that convey information to the market. Volume, timing, orders, traders' identities, related market movements, all have been shown to be important factors in understanding the movement of markets. Indeed, it is likely that information we have not yet even considered may be fundamental in explaining market behavior. This can only be determined by the development of new, ever more complete data sets, leading us back to the question of how to structure data sets that optimize our ability to analyze the rich mosaic of information underlying market behavior.

References

Abhyankar, A., Copeland, L.S., Wong, W., 1995. Nonlinear dynamics in real-time equity market indices: Evidence from the U.K. Economic Journal, forthcoming.

Admati, A.R., Pfleiderer, P., 1988. A theory of intraday patterns: Volume and price variability. Review of Financial Studies 1, 3–40.

Admati, A.R., Pfleiderer, P., 1989. Divide and conquer: A theory of intraday and day-of-the-week effects. Review of Financial Studies 2, 189–223.

Aggarwal, R., Park, Y.S., 1994. The relationship between daily U.S. and Japanese equity prices: Evidence from spot versus futures markets. Journal of Banking and Finance 18, 757–774.

Amihud, Y., Mendelson, H., 1980. Dealership market: Market-making with inventory. Journal of Financial Economics 8, 31–53.

---

[36] The sheer size of the resulting data base may overwhelm all but the most sophisticated computer systems. The scale of the ISSM data (the equity-based US transactions data), for example, is now over 20 Gbyte per year of data, making even storage of the data time-series difficult, to say nothing of its manipulation and analysis. While extensive, this data does not include information on trader identities, the composition of executed trades, or existing limit orders. To the extent that knowledge is facilitated by the multiplicity of research efforts, designing data sets to be complete, but manageable, is a nontrivial task.

Amihud, Y., Mendelson, H., 1987. Trading mechanisms and stock returns: An empirical investigations. Journal of Finance 42, 533–553.

Amihud, Y., Mendelson, H., 1988. Liquidity and asset prices: Financial management implications. Financial Management 17, 5–15.

Andersen, T.G., Bollerslev, T., 1994. Intraday Seasonality and Volatility Persistence in Foreign Exchange and Equity Markets. Kellogg Graduate School of Management, Northwestern University, Working Paper # 193.

Baillie, R.T., 1994. Long Memory Processes and Fractional Integration in Econometrics. Michigan State University, Department of Economics, Unpublished paper, October. Journal of Econometrics, forthcoming.

Baillie, R.T., Bollerslev, T., 1990a. Intra day and inter market volatility in foreign exchange rates. Review of Economic Studies 58, 565–585.

Baillie, R.T., Bollerslev, T., 1990b. A multivariate generalized ARCH approach to modeling risk premia in forward foreign exchange rate markets. Journal of International Money and Finance 9, 309–324.

Baillie, R.T., Bollerslev, T., 1993. Cointegration, fractional cointegration and exchange rate dynamics. Journal of Finance 49, 737–745.

Baillie, R.T., Bollerslev, T., 1994. The long memory of the forward premium. Journal of International Money and Finance 13, 565–571.

Baillie, R.T., Bollerslev, T., Mikkelsen, H.O., 1993. Fractionally integrated generalized autoregressive conditional heteroskedasticity, Journal of Econometrics, forthcoming.

Bank of Japan, 1995. Empirical Analyses of the Information Content of Implied Volatility. Bank of Japan Quarterly Bulletin 3, 64–88.

Barclay, M.J., Warner, J.B., 1993. Stealth trading and volatility: Which trades move prices. Journal of Financial Economics 34, 281–306.

Barclay, M.J., Litzenberger, R.H., Warner, J.B., 1990. Private information, trading volume and stock-return variances. Review of Financial Studies 3, 233–253.

Becker, K.G., Finnerty, J.E., Gupta, M., 1990. The intertemporal relation between the U.S. and the Japanese stock markets. Journal of Finance 45, 1297–1306.

Bennet, P., Kelleher, J., 1988. The international transmission of stock price disruption in October 1987. FRB New York Quarterly Review 13, 17–33.

Benveniste, L.M., Marcus, A.J., Wilhelm, W.J., 1992. What's special about the specialist? Floor exchange versus computerized market mechanisms. Journal of Financial Economics 32, 61–86.

Bessembinder, H., 1994. Bid–ask spreads in the interbank foreign exchange market. Journal of Financial Economics 35, 317–348.

Bhattacharya, S., Constantinides, G.M. (Eds.), 1989. Financial Markets and Incomplete Information: Frontiers of Modern Financial Theory. Rowman and Littlefield, Totowa, NJ.

Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse, Journal of Finance 50, 1655–1690.

Black, F., 1975. Fact and fantasy in the use of options. Financial Analysts Journal 31, 36–72.

Black, F., 1986. Noise. Journal of Finance 41, 529–543.

Block, S.B., French, D.W., McInish, T.H., 1994. NYSE execution costs: A case study of a leading bank. Journal of Financial Services Research 8, 163–175.

Blume, L.E., Easley, D., O'Hara, M., 1994. Market statistics and technical analysis: The role of volume. Journal of Finance 49, 153–181.

Board, J.L.G., Sutcliffe, C.M.S., 1995. The Effects of Trade Transparency in the London Stock Exchange: A Summary. London School of Economics, Financial Markets Group, Discussion Paper.

Bollerslev, T., Domowitz, I., 1992. Some effects of restricting the electronic order book in an automated trade execution system. In: Friedman D. et al. (Eds.), Double Auction Markets: Theories, Institutions, and Experimental Evaluations. Addison-Wesley, CA.

Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. Journal of Finance 48, 1421–1443.

Bollerslev, T., Melvin, M., 1994. Bid–ask spreads and volatility in the foreign exchange market. Journal of International Economics 36, 355–372.

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modelling in finance: A review of the theory and empirical evidence. Journal of Econometrics 8, 5–59.

Bollerslev, T., Domowitz, I., Wang, J., 1994. Order Flow and the Bid-Ask Spread: An Empirical Probability Model of Screen-Based Trading. Working paper, Northwestern University.

Breedon, F., 1993. Intraday price formation on the London Stock Exchange. London School of Economics, Financial Markets Group, Discussion Paper No. 158, March.

Brock, W.A., 1991. Causality, chaos, explanation and prediction in economic and finance. In: Casti, J., Karlqvist, A. (Eds.), Beyond Belief: Randomness, Prediction and Explanation in Science. CRC Press, Boca Raton, ch. 10.

Brock, W.A., Kleidon, A.W., 1992. Periodic market closure and trading volume. Journal of Economic Dynamics and Control 16, 451–489.

Brock, W.A., Le Baron, B., 1993. Using structural modeling in building statistical models of volatility and volume of stock market returns. Department of Economics, University of Wisconsin, Working Paper, December.

Brock, W.A., Lakonishok, J., Le Baron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. Journal of Finance 47, 1731–1764.

Campbell, J.Y., Grossman, S.J., Wang, J., 1993. Trading volume and serial correlation in stock returns. Quarterly Journal of Economics 108, 905–934.

Canina, L., Figlewski, S., 1993. The informational content of implied volatility. Review of Financial Studies 6, 659–681.

Chan, K., 1992. A further analysis of the lead-lag relationship between the cash market and stock index futures market. Review of Financial Studies 5, 123–152.

Chan, K., Chan, K.C., Karolyi, G.A., 1991. Intraday volatility in the stock index and stock index futures market. Review of Financial Studies 4, 657–684.

Chan, K., Chung, Y.P., Johnson, H., 1993. Why option prices lag stock prices: A trading-based explanation. Journal of Finance 48, 1957–1967.

Chang, E.C., P.R. Locke, 1996. The performance and market impact of dual trading: Re CME rule 552. Journal of Financial Intermediation 5, 23–48.

Cheng, M., Madhavan, A., 1994. In Search of Liquidity: Block Trades in the Upstairs and Downstairs Markets. NYSE Working Paper, #94-02.

Christie, A.A., 1981. On Efficient Estimation and Intra-Week Behavior of Common Stock Variances. University of Rochester, unpublished working paper, April.

Clark, P.K., 1973. A subordinated stochastic process model with finite variance for speculative prices. Econometrica 41, 135–155.

Cornell, B., Sirri, E.R., 1992. The reaction of investors and stock prices to insider trading. Journal of Finance 47, 1031–1059.

Dacorogna, M.M., Muller, U.A., Nagler, R.J., Olsen, R., Pictet, O.V., 1993. A geographical model for the daily and weekly seasonal volatility in the FX market. Journal of International Money and Finance 12, 413–438.

de Jong, F., Nijman, T., Roell, A., 1993. A Comparison of the Cost of Trading French Shares on the Paris Bourse and on SEAQ International. London School of Economics Financial Markets Group, Discussion Paper No. 169, August.

Demos, A., Goodhart, C.A.E., 1992. The Interaction between the Frequency of Market Quotations, Spread, and Volatility in the Foreign Exchange Market. LSE Financial Markets Group Discussion Paper, No. 152, October.

Diamond, D.W., Verrecchia, R.E., 1987. Constraints on short-selling and asset price adjustment to private information. Journal of Financial Economics 18, 277–311.

Ding, Z., Granger, C.W.J., Engle, R.F., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.

Domowitz, I., 1993. A taxonomy of automated trading systems. Journal of International Money and Finance 12, 607–631.

Domowitz, I., Wang, J., 1994. Auctions as algorithms. Journal of Economic Dynamics and Control 18, 29–60.

Dutkowsky, D.H., 1993. Dynamic implicit cost and discount window borrowing: An empirical investigation. Journal of Monetary Economics 32, 105–120.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 113–138.

Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. Journal of Finance 47, 577–605.

Easley, D., Kiefer, N.M., O'Hara, M., 1993. One Day in the Life of a Very Common Stock. Working Paper, Cornell University.

Easley, D., O'Hara, M., Srinivas, P.S., 1994. Option Volume and Stock Prices: Evidence on Where Informed Traders Trade. Working Paper, Cornell University.

Easley, D., Kiefer, N.M., O'Hara, M., 1995. The Information Content of the Trading Process. Journal of Empirical Finance, this volume.

Easley, D., Kiefer, N.M., O'Hara, M., Paperman, J., 1996. Liquidity, Information, and Infrequently Traded Stocks. Journal of Finance 51, 1405–1436.

Ederington, L.H., Lee, J.H., 1993. How markets process information: News releases and volatility. Journal of Finance 48, 1161–1191.

Engle, R.F., Russell, J.R., 1995a. Autoregressive Conditional Duration: A New Model for Irregularly Spaced Time Series Data. Working Paper, University of California, San Diego.

Engle, R.F., Russell, J.R., 1995b. Forecast the Frequency of Changes in the Quoted Foreign Exchange Prices with the Autoregressive Conditional Duration Model. Working Paper, University of California, San Diego.

Engle, R.F., Ito, T., Lin, W.-L., 1990. Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market. Econometrica 58, 525–542.

Epps, T.M., Epps, M.L., 1976. The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis. Econometrica 44, 305–321.

Fama, E.F., 1968. The behavior of stock market prices. Journal of Business 38, 34–105.

Finnerty, J.E., Park, H.Y., 1987. Stock index futures: Does the tail wag the dog? A technical note. Financial Analysts Journal 43, 57–61.

Fishman, M.J., Longstaff, F.A., 1992. Dual trading in futures markets. Journal of Finance 47, 643–672.

Foster, F.D., Viswanthan, S., 1993. Can Speculative Trading Explain the Volume-Volatility Relation? Working Paper, Duke University.

Foster, F.D., Viswanathan, S., 1990. A theory of interday variations in volumes, variances, and trading costs in securities markets. Review of Financial Studies 3, 593–624.

Foster, F.D., Viswanathan, S., 1993. Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models. Journal of Finance 48, 187–211.

French, K.R., Roll, R., 1986. Stock return variances: The arrival of information and the reaction of traders. Journal of Financial Economics 17, 5–26.

Gallant, A.R., Rossi, P.E., Tauchen, G., 1992. Stock prices and volume. Review of Financial Studies 5, 199–242.

Ghysels, E., Gourieroux, C., Jasiak, J., 1995. Trading Patterns, Time Deformation and Stochastic Volatility in Foreign Exchange Markets. Working Paper, Universite de Montreal.

Glosten, L.R., 1989. Insider trading, liquidity and the role of the monopolist specialist. Journal of Business 54, 579–596.

Glosten, L.R., 1994. Is the electronic open limit order book inevitable?. Journal of Finance 49, 1127–1162.

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Goodfriend, M.J., 1983. Discount window borrowing, monetary policy, and the post-October 6, 1979 Federal Reserve operating procedure. Journal of Monetary Economics 12, 343–356.

Goodfriend, M.J., Whelpley, W., 1986. Federal funds: Instrument of Federal Reserve policy. Federal Reserve Bank of Richmond Economic Review 72, 3–11.

Goodhart, C.A.E., 1989. News and the Foreign Exchange Market. Manchester Statistical Society, pamphlet, October.

Goodhart, C.A.E., Curcio, R.J., 1992. When Support/Resistance Levels are Broken, can Profits be made? Evidence from the Foreign Exchange Market. London School of Economics, Financial Markets Group, Discussion Paper No. 142, July.

Goodhart, C.A.E., Figliuoli, L., 1991. Every minute counts in financial markets. Journal of International Money and Finance 10, 23–52.

Goodhart, C.A.E., Figliuoli, L., 1992. The geographical location of the foreign exchange market: A test of an islands hypothesis. Journal of International and Comparative Economics 1, 13–28.

Goodhart, C.A.E., Giugale, M., 1993. From hour to hour in the foreign exchange market. The Manchester School 61, 1–34.

Goodhart, C.A.E., Payne, R.J., 1995. Microstructural Dynamics in a Foreign Exchange Electronic Broking System. London School of Economics, Financial Markets Group, February, unpublished paper.

Goodhart, C.A.E., Hall, S.G., Henry, S.G.B., Pesaran, B., 1993. News effects in a high frequency model of the sterling–dollar exchange rate. Journal of Applied Econometrics 8, 1–13.

Goodhart, C.A.E., Ito, T., Payne, R.J., 1994. One day in June, 1993: A Study of the Working of Reuters 2000-2 Electronic Foreign Exchange Trading System. Paper presented at the NBER Conf. on The Microstructure of Foreign Exchange Markets, Perugia, Italy, July. Frankel, J., Galli, G., Giovannini, A. (Eds.). Proceedings forthcoming.

Granger, C.W.J., Morgenstern, O., 1970. Predictability of Stock Market Prices. Heath-Lexington, Lexington, MA.

Guillaume, D., Dacorogna, M.M., Dave, R., Muller, U.A., Olsen, R.B., Pictet, O.V., 1994. From the Bird's Eye to the Microscope: A Survey of new Stylized Facts of the Intra-Daily Foreign Exchange Markets. A discussion paper by the O and A research group, Zurich, October.

Guillaume, D., Pictet, O.V., Dacorogna, M.M., 1995. On the Intraday Performance of GARCH Processes. Working Paper, Olsen and Associates, Zurich, Switzerland.

Hamao, Y., Masulis, R.W., Ng, V., 1990. Correlations in price changes and volatility across international stock markets. Review of Financial Studies 3, 281–308.

Hamao, Y., Masulis, R.W., Ng, V., 1991. The effect of the 1987 stock crash on international financial integration. In: Zeimba, W., Bailey, W., Hamao, Y. (Eds.), Japanese Financial Market Research. Elsevier Science Publishers, pp. 483–502.

Han, S., Kalay, A., Rosenfeld, A., 1994. The Information Content of Non-Trading. Working paper, University of Utah.

Harris, L., 1986. A transactions data study of weekly and intradaily patterns in stock returns. Journal of Financial Economics 16, 99–117.

Harris, L., 1989. The October 1987 S&P 500 Stock-Futures Bases. Journal of Finance 44, 77–99.

Harvey, A.C., Shephard, N., 1993. The Econometrics of Stochastic Volatility. London School of Economics, Financial Markets Group, Discussion Paper No. 166.

Harvey, C.P., Whaley, R.E., 1992. Market volatility prediction and the efficiency of the S&P 100 index option market. Journal of Financial Economics 31, 43–73.

Harvey, A.C., Ruiz, E., Shephard, N., 1994. Multivariate stochastic variance models. Review of Economic Studies 61, 247–264.

Hasbrouck, J., 1988. Trades, quotes, inventories and information. Journal of Financial Economics 22, 229–252.

Hasbrouck, J., 1991a. Measuring the information content of stock trades. Journal of Finance 46, 179–207.

Hasbrouck, J., 1991b. The summary informativeness of stock trades: An econometric analysis. Review of Financial Studies 4, 571–595.

Hasbrouck, J., Ho, T.S.Y., 1987. Order arrival, quote behavior and the return generating process. Journal of Finance 42, 1035–1048.

Hasbrouck, J., Sofianos, G., 1993. The trades of market makers: An empirical analysis of NYSE specialists. Journal of Finance 48, 1565–1593.

Hasbrouck, J., Sofianos, G., Sosebee, D., 1993. The New York Stock Exchange Systems and Trading Procedures. NYSE Working Paper, No. 93-01.

Hausman, J.A., Lo, A.W., 1990. A Continuous Time Discrete State Stochastic Process for Transactions Stock Prices: Theory and Empirical Evidence. Working paper, MIT.

Ho, T.S.Y., Stoll, H.R., 1983. The dynamics of dealer markets under competition. Journal of Finance 38, 1053–1075.

Hogan, K., Melvin, M., 1994. Sources of meteor showers and heat waves in the foreign exchange market. Journal of International Economics 37, 239–247.

Holthausen, R.W., Leftwich, R.W., Mayers, D., 1987. The effect of large block transactions on security prices: A cross-sectional analysis. Journal of Financial Economics 19, 237–267.

Huang, R.D., Stoll, H.R., 1994. Market microstructure and stock market predictions. Review of Financial Studies 7, 179–213.

Hutchinson, J.M., Lo, A.W., Poggio, T., 1994. A nonparametric approach to pricing and hedging derivative securities via learning networks. Journal of Finance 49, 851–889.

Ito, T., Engle, R.F., Lin, W.-L., 1992. Where does the meteor shower come from?. Journal of International Economics 32, 221–240.

Jacquier, E., Polson, N.G., Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. Journal of Business and Economic Statistics 12, 371–417.

Jaffe, J.F., 1974. Special information and insider trading. Journal of Business 47, 410–428.

Jain, P.C., Joh, G-H., 1988. The dependence between hourly prices and trading volume. Journal of Financial and Quantitative Analysis 23, 269–283.

Jones, C.M., Kaul, G., Lipson, M.L., 1994. Transactions, volume, and volatility. Review of Financial Studies 7, 631–651.

Jorion, P., 1994. Risk and Turnover in the Foreign Exchange Market. Paper presented at NBER Conf. on The Microstructure of Foreign Exchange Markets, Perugia, July. Frankel, J., Galli, G., Giovannini, A. (Eds.). Forthcoming as proceedings.

Karpoff, J.M., 1987. The relation between price changes and trading volume: A survey. Journal of Financial and Quantitative Analysis 22, 109–126.

Kawaller, I., Koch, P., Koch, T., 1987. The temporal price relationship between S&P 500 futures and the S&P 500 futures index. Journal of Finance 42, 1309–1329.

Kawaller, I.G., Koch, P.D., Koch, T.W., 1990. Intraday relationships between volatility in S&P 500 futures prices and volatility in the S&P 500 index. Journal of Banking and Finance 14, 373–397.

Keim, D.B., Madhavan, A., 1996. The upstairs market for large-block transactions: Analysis and measurement of price effects. Review of Financial Studies 9, 1–36.

Kim, O., Verrechia, R.E., 1991. Market reactions to anticipated announcements, Journal of Financial Economics 30, 273–310.

King, M.A., Wadhwani, S., 1990. Transmission of volatility between stock markets. Review of Financial Studies 3, 5–33.

King, M.A., Sentana, E., Wadhwani, S., 1990. A Heteroscedastic Factor Model of Asset Returns and Risk Premia. London School of Economics, Financial Markets Group, Discussion Paper No. 80, May.

King, M.A., Sentana, E., Wadhwani, S., 1993. Volatility and Links between National Stock Markets. London School of Economics, Financial Markets Group, Discussion Paper No. 173, October.

Kleidon, A.W., Werner, I.M., 1994. Round-the-clock Trading: Evidence from UK Cross-Listed Securities. London School of Economics, Financial Markets Group, Discussion Paper No. 182, February.

Kofman, P., Bouwman, A., Moser, J., 1994. Is there LIF(F)E after DTB: Price Dynamics for Cross Listed Futures Contracts on Synchronous Markets. Federal Reserve Bank of Chicago, Research Department, mimeo.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.

Laffont, J.-J., Maskin, E., 1990. The efficient market hypothesis and insider trading on the stock market. Journal of Political Economy 98, 70–93.

Lamoureux, C.G., Lastrapes, W.D., 1990. Heteroskedasticity in stock return data: Volume versus GARCH effects. Journal of Finance 45, 221–228.

Laux, P.A., Ng, L.K., 1993. The sources of GARCH: Empirical evidence from an intraday returns model incorporating systematic and unique risks. Journal of International Money and Finance 12, 543–560.

Le Baron, B., 1992. Do Moving Average Trading Rule Results imply Nonlinearities in Foreign Exchange Markets? University of Wisconsin, technical report.

Le Baron, B., 1994. Technical Trading Rule Profitability and Foreign Exchange Intervention. University of Wisconsin, Department of Economics, SSRI paper 9445, October.

Lee, C.M.C., Ready, M.J., 1991. Inferring trade direction from intraday data. Journal of Finance 46, 733–746.

Lee, C.M.C., Shleifer, A., Thaler, R.H., 1991. Investor sentiment and the closed end fund puzzle. Journal of Finance 46, 75–110.

Lee, C.M.C., Mucklow, B., Ready, M.J., 1993. Spreads, depth, and the impact of earnings information: An intraday analysis. Review of Financial Studies 6, 345–374.

Lehmann, B.N., Modest, D.M., 1994. Trading and Liquidity on the Tokyo Stock Exchange: A Birds Eye View. University of California, San Diego, Finance Working Paper No. 234, April.

Levich, R.M., Thomas, L.R., 1993. The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach. Journal of International Money and Finance 12, 451–474.

Lin, W.-L., 1994. Interdependence Between International Stock Returns – News or Contagion Effect? University of Wisconsin, SSRI Working Paper 9442, September.

Lin, W.-L., Engle, R.F., Ito, T., 1991. Do Bulls and Bears move across Borders? International Transmission of Stock Returns and Volatility as the World Turns. NBER working paper, No. 3911.

Lo, A.W., 1994. Neural Networks and Other Nonparametric Techniques in Economics and Finance, Blending Quantitative and Traditional Equity Analysis. Association for Investment Management and Research, pp. 25–36.

Lo, A.W., MacKinley, A.C., 1988. Stock market prices do not follow random walks: Evidence from a simple specification test. Review of Financial Studies 1, 3–40.

Lockwood, L.J., Linn, S.C., 1990. An examination of stock market return volatility during overnight and intraday periods, 1964–1989. Journal of Finance 45, 591–601.

Lyons, R.K., 1994. Foreign Exchange Volume: Sound and fury Signifying Nothing. Paper presented at the NBER Conf. on Microstructure of Foreign Exchange Markets, Perugia, Italy, July. Frankel, J., Galli, G., Giovannini, A. (Eds.). Proceedings forthcoming.

Lyons, R.K., 1995. Tests of microstructural hypotheses in the foreign exchange market. Journal of Financial Economics 39, 321–351.

Lyons, R.K., 1996. Optimal transparency in a dealer market with an application to foreign exchange, Journal of Financial Intermediation 5, 225–254.

Madhavan, A., 1992. Trading mechanisms in securities markets. Journal of Finance 47, 607–641.

Madhavan, A., Smidt, S., 1991. A Bayesian model of intraday specialist pricing. Journal of Financial Economics 30, 99–134.

Madhavan, A., Smidt, S., 1993. An analysis of changes in specialist inventories and quotations. Journal of Finance 48, 1595–1628.

Madhavan, A., Richardson, M., Roomans, M., 1994. Why do Security Prices Change? A Transactions-Level Analysis of NYSE Stocks. University of Southern California, unpublished paper, July.

Manaster, S., Rendleman, R.J., 1982. Option prices as predictors of equilibrium stock prices. Journal of Finance 37, 1043–1058.

Mandelbrot, B.B., Taylor, H.M., 1967. On the distribution of stock price differences. Operations Research 15, 1057–1062.

McInish, T.H., Wood, R.A., 1990a. A transaction data analysis of the variability of common stock returns during 1980–1984. Journal of Banking and Finance 14, 99–112.

McInish, T.H., Wood, R.A., 1990b. An analysis of transactions data for the Toronto Stock Exchange. Journal of Banking and Finance 14, 458–491.

McInish, T.H., Wood, R.A., 1991. Hourly returns, volume, trade size and number of trades. Journal of Financial Research 14, 303–315.

McInish, T.H., Wood, R.A., 1992. An analysis of intraday patterns in bid/ask spreads for NYSE stocks. Journal of Finance 47, 753–764.

Menkhoff, L., 1995. Examining the Use of Technical Currency Analysis. University of Freiburg, Department of Economics, Working Paper, January.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous-time model. Journal of Economic Theory 3, 373–413.

Meulbrock, L.K., 1992. An empirical analysis of illegal insider trading. Journal of Finance 47, 1661–1699.

Muller, U.A., Sgier, R., 1992. Statistical Analysis of Intra-Day Bid–Ask Spreads in the Foreign Exchange Market. Internal document UAM 04-10, Olsen and Associates, Zurich.

Muller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M., Morgenegg, C., 1990. Statistical study of foreign exchange rates, empirical evidence of a price scaling law, and intraday analysis. Journal of Banking and Finance 14, 1189–1208.

Nelson, D.B., 1990. ARCH models as diffusion approximations. Journal of Econometrics 45, 7–38.

Neuberger, A., 1992. An empirical examination of market maker profits on the London Stock Exchange. Journal of Financial Services Research 6, 343–372.

Neuberger, A., Roell, A., 1992. Components of the Bid–Ask Spread: A Glosten–Harris Approach. London Business School, working paper.

Neumark, D., Tinsley, P.A., Tosini, S., 1991. After-hours stock prices and post-crash hangovers. Journal of Finance 45, 159–178.

Ng, N., 1987. Detecting spot prices forecasts in futures prices using causality tests. Review of Futures Markets 6, 61–89.

O'Hara, M., 1995. Market Microstructure Theory. Blackwell Publishers, Cambridge, MA.

O'Hara, M., Oldfield, G.S., 1986. The microeconomics of market making. Journal of Financial and Quantitative Analysis 21, 361–376.

Oldfield, G.S., Rogalski, R.J., 1980. A theory of common stock returns over trading and non-trading periods. Journal of Finance 35, 729–751.

Pagano, M., Roell, A., 1990a. Trading Systems in European Stock Exchanges. London School of Economics, Financial Markets Group, Discussion Paper, No. 75, February.

Pagano, M., Roell, A., 1990b. Auction Markets, Dealership Markets and Execution Risk. London School of Economics, Financial Market Group, Discussion Paper, No. 102, October.

Pagano, M., Roell, A., 1991. Dually-Traded Italian Equities. London School of Economics, Financial Markets Group, Discussion Paper, No. 116, April.

Pagano, M., Roell, A., 1992. Auction and dealership markets: What is the difference?. European Economic Review 36, 613–623.

Pagano, M., Roell, A., 1995. Transparency and liquidity: A comparison of auction and dealer markets with informed trading. Journal of Finance, forthcoming.

Petersen, M.A., Fialkowski, D., 1994. Posted versus effective spreads: Good prices or bad quotes?. Journal of Financial Economics 35, 269–292.

Pictet, O.V., Dacorogna, M.M., Muller, U.A., Olsen, R.B., Ward, J., 1992. Real-time trading models for Foreign Exchange rates, Neural Network World 6, 713–744.

Porter, D.C., 1992. The probability of a trade at the ask: An examination of interday and intraday behavior. Journal of Financial and Quantitative Analysis 27, 209–227.

Pulli, M., 1992. Overnight Market Interest Rates and Banks Demand for Reserves in Finland. Bank of Finland Institute for Economic Research, Series B, 47.

Richardson, M., Smith, T., 1994. A direct test of the mixture of distributions hypothesis: Measuring the daily flow of information. Journal of Financial and Quantitative Analysis 29, 101–115.

Roll, R., 1984. A simple implicit measure of the effective bid–ask spread in an efficient market. Journal of Finance 39, 1127–1159.

Schnadt, N., 1994. The Domestic Money Markets of the U.K., France, Germany and the U.S. Subject Report VII, Paper 1, City Research Project, London Business School.

Schwert, G.W., 1990. Stock volatility and the crash of 87. Review of Financial Studies 3, 77–102.

Seppi, D.J., 1992. Block trading and information revelation around quarterly earnings announcements. Review of Financial Studies 5, 281–305.

Sheikh, A.M., Ronn, E.J., 1994. A characterization of the daily and intraday behavior of returns on options. Journal of Finance 49, 557–579.

Smith, T., R. Whaley, 1995. Assessing the costs of dual trading. Journal of Law and Economics, forthcoming.

Snell, A., Tonks, I., 1995. Determinants of price quote revisions on the London Stock Exchange. Economic Journal 105, 77–94.

Stephan, J., Whaley, R.E., 1990. Intraday price change and trading volume relationships in the stock and stock options markets. Journal of Finance 45, 191–220.

Stoll, H.R., Whaley, R.E., 1990. Stock market structure and volatility. Review of Financial Studies 3, 37–71.

Susmel, R., Engle, R.F., 1994. Hourly volatility spillovers between international equity markets. Journal of International Money and Finance 13, 3–25.

Tauchen, G.E., Pitts, M., 1983. The price variability–volume relationship on speculative markets. Econometrica 51, 485–505.

Taylor, M.P., Allen, H., 1992. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11, 304–314.

Trippi, R.R., E. Turban, 1992. Neural Network in Finance and Investing. Probus Publishing, Chicago, IL.

Vassilicos, C., 1990. Are Financial Markets Chaotic? A Preliminary Study of the Foreign Exchange Market. London School of Economics, Financial Markets Group, Discussion Paper No. 86, June.

Vassilicos, C., Demos, A., 1994. The Multifractal Structure of High Frequency Foreign Exchange Rate Fluctuations. London School of Economics, Financial Markets Group, Discussion Paper No. 195, September.

Vassilicos, C., Demos, A., Tata, F., 1992. No Evidence of Chaos but some Evidence of Multifractals in the Foreign Exchange and the Stock Markets. London School of Economics, Financial Markets Group, Discussion Paper No. 143, June.

Vijh, A.M., 1988. Potential biases from using only trade prices of related securities on different exchanges: A comment. Journal of Finance 43, 1049–1055.

Vila, A.F., Sandmann, G., 1994. Floor Trading versus Electronic Screen Trading: An Empirical Analysis of Market Liquidity and Information Transmission in the Nikkei Stock Index Futures Markets. Boston University, unpublished paper, May.

Von Furstenberg, G.M., Jeon, B.N., 1989. International stock price movements: Links and messages. Brookings Papers on Economic Activity 1, 125–167.

Wood, R.A., McInish, T.H., Ord, J.K., 1985. An investigation of transaction data for NYSE stocks. Journal of Finance 40, 723–741.

Zabel, E., 1981. Competitive price adjustment without market clearing. Econometrica 49, 1201–1221.