

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: The Microstructure of Foreign Exchange Markets

Volume Author/Editor: Jeffrey A. Frankel, Giampaolo Galli, Alberto Giovannini, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-26000-3

Volume URL: <http://www.nber.org/books/fran96-1>

Conference Date: July 1-2, 1994

Publication Date: January 1996

Chapter Title: One Day in June 1993: A Study of the Working of the Reuters 2000-2 Electronic Foreign Exchange Trading System

Chapter Author: Charles Goodhart, Takatoshi Ito, Richard Payne

Chapter URL: <http://www.nber.org/chapters/c11364>

Chapter pages in book: (p. 107 - 182)

4 One Day in June 1993: A Study of the Working of the Reuters 2000-2 Electronic Foreign Exchange Trading System

Charles Goodhart, Takatoshi Ito, and Richard Payne

4.1 Introduction

This is a study of foreign exchange dealers' behavior as revealed in the working of Reuters 2000-2, a recently developed electronic foreign exchange trading system. It was launched in 1992 with twenty-three subscriber sites in two countries and by September 1993 had more than 230 dealing sites in twenty-eight cities in seventeen countries (Blitz 1993). The working of the system is described in more detail in section 4.2. This dealing system 2000-2 (henceforward termed D2000-2) is, however, still at the developing rather than a mature stage, and the snapshot that we have of its operations on one day—

Charles Goodhart is the Norman Sosnow Professor of Banking and Finance and deputy director of the Financial Markets Group at the London School of Economics. Takatoshi Ito is professor of economics at Hitotsubashi University and senior advisor of the Research Department at the International Monetary Fund. Richard Payne is a Ph.D. student at the London School of Economics and a research assistant at the Financial Markets Group.

This lengthy empirical exercise was conducted in a number of stages. After one of the authors, C. Goodhart, had obtained the original videotapes from Reuters, to whom we are most grateful, the data on the tapes were transcribed onto paper by two of the authors' wives, Mrs. Goodhart and Mrs. Ito, assisted by Yoko Miyao, a painstaking task beyond and above the normal requirements of matrimony. The data were then sorted and organized by T. Ito and R. Payne, separately in the United States and the United Kingdom. The graphic appendix is entirely Ito's work. The descriptive material in sections 4.1 and 4.2 was mostly written by Goodhart. The comparison of D2000-2 and FXFX in section 4.3 had input from all authors, but mostly Goodhart and Payne. The comparable FXFX data were obtained from Olsen and Associates, to whom we are most grateful. Only the first three sections were ready in time for the July Perugia conference, so this is all that our discussants, to whom we are most grateful, then had before them. Section 4.4, completed thereafter, was entirely the work of Goodhart and Payne, with Payne responsible for the econometrics, apart from table 4.16 by Ito. Charles Goodhart and Richard Payne wish to thank the Economic and Social Research Council for financial support. Takatoshi Ito thanks Charles Kramer for technical assistance in producing the graphic appendix.

16 June 1993—may have become outdated and obsolete by the time that this is published.¹

Reuters has become subject to competition in this marketplace, from Minex and from the Electronic Broking Service (EBS). The former was established in April 1993 by Japanese institutions and, according to Blitz (1993), is “much used in Asia,” although, as of September 1993, it did not reveal the number of trades crossed or terminals used. EBS was founded on Wednesday, 21 September 1993. It cost, again according to Blitz, around £40 million to launch and has been backed by a dozen leading banks in foreign exchange—such as Citibank and Chase Manhattan—who formed a consortium with Quotron, an electronic information screen competitor with Reuters.

In September 1993, Bob Etherington, Reuters’ international marketing manager, would not reveal his dealing system’s current volume levels, although Blitz (1993) did report that the “system has reached [its] initial target of 1,000 trades a day, each for a minimum 1 million units of currency dealt.”² As noted, Minex was not then disclosing the number of trades, and EBS had not started but was going to invite dealers “to trade in standard amounts of \$5 million in Dm/\$ and £5 million in £/Dm.”

Such electronic *dealing* systems (as contrasted with informational pages supplying indicative bid-ask quotes, such as the Reuters FXFX page) are still in their early stages and are highly competitive. Moreover, they may have an important future: “Roughly 60 per cent of deals in the currency market are now done by traders in two banks—or counterparties—who call one another up directly. The remainder of deals are done through brokers, who bring together diverse buyers and sellers. . . . But they [the banks] complain that the commissions charged for broking a deal are very high. Automated brokerage terminals do the same job as humans at a reduced cost. . . . The banks are attracted by the reduced cost of commission. But they fear that 2000-2 will help monopolize the market in electronic dealing systems. Mr. Bartko [chairman of the EBS partnership] admits that this is one of the principal motives for this week’s launch of EBS” (Blitz 1993).

Electronic trading systems have been in use for rather longer in other financial markets, notably in standardized futures and options markets. Instinet and Globex are two such that Reuters has again been developing. A useful taxonomy of the modus operandi of such electronic trading systems has been provided by Domowitz (1990, 1993).

1. Readers wanting more up-to-date information should refer directly to Reuters Limited, 85 Fleet Street, London EC4P 4AJ, United Kingdom.

2. The total amount thus traded is large in absolute amount but small relative to reported daily turnover in this market of some \$900 billion or more. We find it hard to relate the data reported above to the BIS (1993) report in their 1992 survey that, “in the United States and the United Kingdom, the share of deals going through such [automated dealing] systems in April 1992 was 32 and 24% respectively” (table 1, p. 21, and p. 24). Probably definitions of automated dealing systems would have been somewhat wider, including Reuters D2000-1 as well as D2000-2, but, even so, the above percentage seems surprisingly high.

Under these circumstances, details of the workings of such systems remain commercially sensitive. The database that we have studied, a videotape of all the entries over D2000-2 for almost exactly seven hours for the deutsche mark/dollar, and some sixteen minutes less for five other bilateral exchange rates, shown on the D2000-2 screen during European business hours on 16 June 1993 (from 08:31:50 to 15:30:00 British Standard Time [BST], i.e., GMT + 1), remains the copyright of Reuters.³ Anyone wishing to use these data should refer to Reuters, not to us. We should like to emphasize that this videotape did not include, and we have not been given any access to, any information regarding the identity of any of the parties involved in trading; all the trades observed by us remain anonymous. Indeed, it is not possible for any observer, even in Reuters itself, to identify which are the individual banks using the system.

Readers should keep in mind the shortcomings of these data. They represent a short snapshot of conditions in a rapidly changing market over a year ago. Trading undertaken over such electronic trading systems may well be, as discussed further below, not representative of the market as a whole; trading activity on D2000-2 on 16 June 1993 may have differed in some respects significantly from that in surrounding days and weeks; the volume and characteristics of electronic trading (over Reuters) in June 1993 may well be quite different from that now since over a year has passed.

Given these disclaimers, why should anyone bother to read on? Despite these shortcomings, there are, however, several reasons why this study provides new insights in the literature of high-frequency exchange rate behavior. First, until now there have been virtually no *continuous* time-series data available at all on actual trades, prices, and volumes in the foreign exchange market.⁴ The 60 percent or so of deals done directly by two bank counterparties over the telephone remain, naturally, private information. There has been little use made of data on foreign exchange *transactions* intermediated by specialist interbank brokers, no doubt partly because of commercial and confidentiality sensitivities. The only studies currently known to us making use of such data are by Lyons (1995, chap. 5 in this volume). Data of any kind on the characteristics and continuous time-series behavior of actual trading transactions on the foreign exchange market are, therefore, still rare.⁵ Second, there have been so few data on *transactions* in the foreign exchange market that almost all the

3. We are most grateful to Reuters in general and to Mr. Etherington in particular for allowing us to record the quantitative details reported below.

4. There is, of course, the survey of foreign exchange business that has now been undertaken three times at three-year intervals in April 1986, 1989, and 1992 by central banks under the aegis of the Bank for International Settlements (BIS), but this does not provide time-series data. The volumes reported are aggregates for the month of April.

5. We have little doubt that such data will become more plentiful and easily available in the future. But for the time being at least they have rarity value. Also, as electronic trading systems mature, it should be of historical interest to observe how they looked and operated in the early stages of their development.

studies on this market have used data on bilateral currency exchange rates that emanate from the *indicative* bid-ask prices shown on electronic screens by the specialist information providers, for example, Reuters, Telerate, Knight Ridder, and Quotron. There has, naturally, been some concern whether the high-frequency characteristics of such indicative quotes, for example, the negative auto-correlation and the fact that the size of the spread clusters at certain conventional values, are representative of the characteristics of *firm* (committed) bid-ask quotes at the touch. The *touch*, a term more commonly used in the United Kingdom than in the United States, is defined as the difference between the best (highest) bid and the lowest ask on offer, where these are (usually) input by different banks. Lyons, for example, expressed such concerns when he wrote, "Some of the shortcomings of the indicative quotes include the following. First, they are not transactable prices. Second, while it is true that the indicated spreads usually bracket actual quoted spreads in the interbank market, they are typically two to three times as wide. . . . Third, the indications are less likely to bracket true spreads when volatility is highest since there are limits to how frequently the indications can change. And finally, my experience sitting next to dealers at major banks indicates that they pay no attention at all to the current indication; rather, dealers garner most of their high-frequency market information from signals transmitted via intercoms connected to inter-dealer brokers [see Lyons 1993]. In reality, the main purpose of the indicative quotes is to provide non dealer participants with a gauge of where the inter-dealer market is trading" (1995, pp. 331–32; see also Flood 1994, esp. n. 6, p. 154).

Do, for example, the frequency and volatility of the *indicative* quotes provide a reasonable proxy for the same characteristics both in the *committed* bid-ask quotes and in the associated *transactions* in the electronic trading systems? We provide an initial answer to such questions in section 4.3, where we seek to compare characteristics of the FXFX time series⁶ with those of the D2000-2 data for the overlapping seven hours. As described in more detail in section 4.3, the D2000-2 series was *not* time-stamped, and our study of this relation is conditional on the assumptions and techniques used to match these two series temporally.

Subject to that condition, and to anticipate some of our main findings in section 4.3, the averages of the bid-ask in both series (FXFX and D2000-2) are almost identical. A graph of the time path for the deutsche mark/dollar from the two sources looks like one line (see figure 4.1). Thus, the time path of the indicative quotes *can*, on this evidence, be taken as a very good and close proxy for that in the underlying *firm* series. Nevertheless, some of the characteristics of the bid-ask series, for example, the pattern of autocorrelation, are somewhat different. Even so, both series indicate a somewhat similar GARCH

6. We obtained the accompanying FXFX data series from Dr. M. Dacorogna of Olsen and Associates in Zurich.




Fig. 4.1 Average of bid-ask for FFXF and D2000-2 data: deutsche mark/dollar

pattern. As would be expected, the two series are cointegrated, with the indicative series responding more to deviations from the equilibrium (i.e., a larger and more significant negative coefficient on the error correction mechanism). By contrast, the characteristics of the spreads in the FFXF as compared with the touch in D2000-2 are markedly different. The spreads in the FFXF series show clustering among a small number of standard values (e.g., 5, 7, and 10 pips for the deutsche mark/dollar), whereas the spreads at the touch show no such signs of clustering.

After examining the relations between the quote series and associated spreads of FFXF and D2000-2 in section 4.3, we turn in section 4.4 to a more detailed study of the characteristics of D2000-2, in particular, the interaction between quotes and transactions in that data set. This long section has five subsections. First, in section 4.4.1, we examine the statistical characteristics of the transaction price series in D2000-2. Whereas for both D2000-2 and FFXF the quote series incorporate a first-order negative moving average, the transaction price data appear to follow a random walk. Our most interesting finding is that the series of runs of deals, sequences of trades at the bid and the ask, is not normally distributed but contains some very long consecutive sequences, another fat-tailed distribution.

Second, in section 4.4.2, we examine the interrelations between the available data series, using nine main series from D2000-2, all of which, apart from the spread, can be separately obtained for the bid and the ask. These are the frequency of transactions (deals), their size, and whether such transactions exhausted the quantity currently quoted; then the frequency of quote revision, the change in the quoted prices, and the quantity quoted; and two measures of volatility, the absolute change in the quote and the standard deviation of the quotes. Our main finding is that there is a two-way interrelation between the

frequency of quote revisions and the *frequency* of deals and that, when a deal exhausts the quantity on offer, this then affects (with one-way causality) a nexus of relations between volatility, spreads, and quote revisions. We also conduct similar companion studies on the (temporally associated) FAFX data using a smaller subset of data series (since we have no data on transaction characteristics or on posted quantities from FAFX), but these have less interesting results.

Our finding that there is a strong two-way relation between the frequency of quote revisions and that of transactions within a period is, we believe, new, although the underlying cause, that both derive from the arrival of "news," is theoretically straightforward. Most studies of transactions in other asset markets (e.g., the New York Stock Exchange [NYSE]), have used data series calibrated in transaction (tick) time, with the result that one cannot then infer calendar-time frequency. Otherwise, with relatively low-frequency transactions on the NYSE, so many of the observations would exhibit zero change. With much higher-frequency transactions on foreign exchange markets, it seemed to us worthwhile to explore the form of these relations in both clock time and transaction time, although we feel that much remains to be done in clarifying the appropriate econometric usage in this field.

Next, in section 4.4.3, we examine the ARCH (autoregressive conditional heteroskedasticity) characteristics of the quote series, in particular to discover whether their GARCH characteristics would be affected by the addition of transactions data. In this case, unlike most of the other main results in section 4.4.2, the results did appear sensitive to whether the exercise was run in clock time or tick time.

Largely because much more data have been made available for the equity market, especially the NYSE, and its associated derivative markets, there has been much more empirical work on those markets than for the foreign exchange market. Moreover, the two markets are quite dissimilar in format and microstructure, as nicely described in Bessembinder (1994). Nevertheless, despite the comparatively very small size of our data set, its coverage of transactions as well as quotes brings it somewhat nearer to the richer data sets available on equity markets. In particular, our study here, examining the interaction between trades at the bid and ask and price quote revisions, has some features in common with that of Hasbrouck's (1991) study of such effects in the NYSE. So we then replicate his study as closely as we can, using our own data set and adding some variations of our own.

We draw the conclusions of these exercises undertaken earlier in section 4.4 together in the final part, section 4.4.5. Throughout this work, the caveat that our data set lasts for only seven hours, a possibly atypical period, must always be kept in mind, despite the comparatively large number of data points. It is in this sense a very small sample. All our findings, both positive and negative, must be treated with caution.

4.2 The Characteristics of D2000-2

Automated brokerage terminals do the same job as humans but at a reduced cost. A bank dealer who is a member of one of these electronic systems can enter her buy and/or sell price into them. Reuters D2000-2 and EBS show only the touch, the highest bid, and the lowest ask; these will normally, but not necessarily, be entered by different banks. This is different from the indicative foreign exchange pages (e.g., FAFX), which show the *latest* update of the bid and ask entered by a single identified bank. On all the electronic trading systems, the identity of the inputting bank is *not* shown. The quantity that the inputting bank is prepared to trade is also shown on D2000-2. This was then shown as integers of \$1 million, and in some bilateral cases DM 1 million, from 1-5 and entered as M (medium) for a sum between \$6 and \$10 million and L (large) for sums above \$10 million.⁷ More than one bank may input the same best bid (ask) price, in which case the quantity shown is the sum of that offered by these banks. The limit orders, that is, those below the (best) bid and above the (best) ask, and their associated firm quantities are entered and stored in these systems but are not revealed over D2000-2 and EBS. Such reserve limit orders are shown on Minex.

Another bank dealer and member of the trading system can then "hit" either the bid or the ask by typing instructions on his own machine. The first check is prudential. Banks in such systems may want to restrict the amount of dealing with certain other counterparties (in some cases refusing to deal at all with some counterparties). The computer first checks whether the deal is prudentially acceptable to both parties (who remain at this stage anonymous). If not, the deal is refused and the "hitter" so informed. We have no information as to how often this might happen, but we surmise that it might be fairly rare. Assuming that the "hit" is accepted and that several banks are offering the same best price, their offers are met on the basis of the time of entry, first in first out. When a new deal is made, the new transaction price enters on the right-hand column of the screen,⁸ and there must be an associated change in the quantity of the bid (ask), depending on which is hit,⁹ and also in the price offered if the size of the deal exhausts the quantity offered at the previous price. In such cases, the bid price must move downwards if there was an exhaustive deal at the bid, and the ask price upward following an exhaustive deal at the ask, or indicate that there are no remaining limit bids (asks) in the systems, that is, no quote shown.¹⁰ Note that, in an automatic system like this, a deal must be made

7. This classification has since been changed.

8. When a new deal has been made, the new transactions price initially for a few seconds shows purple, rather than the standard black, on the screen in order to alert traders to this.

9. When the deal is completed, both banks, the hitter and the quoter, will be sent details regarding to whom and where to make the payment, which is then settled in the standard fashion. So, ex post facto, the identity of the counterparty becomes revealed.

10. Unhappily, we had a few cases in our data where this directional constraint did not hold. While this could be due to new bid-ask inputs occurring at exactly the same moment, several of

at either the posted bid or ask and cannot be made at an interior price between them, as can happen with nonautomated human dealers, which can cause problems in empirical studies. This has been a particular problem for empirical studies of the NYSE (see, e.g., Petersen and Fialkowski 1994; and Lee and Ready 1991).

D2000-2 allowed traders to deal in some fifteen major bilateral exchange rates at the time of our exercise. The number and range of currencies covered have been changing over time, as is no doubt the case for EBS and Minex as well. The screen for D2000-2 is not big enough to show all fifteen at once, and in any case such a large number of separate rates might be distracting. So the dealer on D2000-2 can call up to six bilateral exchange rate onto the screen at any one time.

All this may be made somewhat easier to follow by seeing an example of what a dealer would see when looking at her screen. This is shown in table 4.1. Note, in particular, that not all the cells have entries. There are periods, especially in the less actively traded bilateral exchange rates, when no bank is making a firm offer. A bilateral currency can have a firm bid (ask) exhibited without there being any corresponding ask (bid) on the screen, as in this example for the deutsche mark/French franc exchange rate; so there is no observed spread at such times. Any bid-ask price must be associated with an accompanying quantity offered (and vice versa). As electronic trading becomes more popular, such gaps in prices may be expected to become fewer. Note also that the representation of the bilateral exchange rate in the left-hand column is the reverse of what would be normally expected, that is, row 1 would in normal usage be described as the number of yen per dollar. (We thank a discussant for noticing this.) The reason, we understand, for this ordering is that all the volumes are denominated in units of 1 million of the *first* currency shown. Henceforth, however, we will revert to the standard representation of the bilateral rates.

D2000-2 runs throughout the whole day during the week, apart from a short break from 2300 GMT to 0100 GMT. On 16 June 1993, a Reuters employee started to videotape the bilateral deutsche mark/dollar exchange rate at approximately 0830 hours BST. This is the dominant and most active of all exchange rates (see, e.g., Goodhart and Demos 1990, 1991a, 1991b). About sixteen minutes, thirty seconds later, he also put the additional five bilateral exchange rates that were shown in table 4.1 up onto the screen.¹¹

these cases probably arise from mistakes in transcribing the videotape (see section 4.2). When we had identified these few errors, we removed them from the data set.

11. Reuters had decided to videotape a day (seven hours) of the working of D2000-2 for their own purposes. We do not know why their operator chose these other five bilateral exchange rates. There is some autocorrelation in volatility and activity in differing rates from day to day, and maybe the operator felt that these would provide either more interest or a better representation than the other nine available. But, basically, we do not know, just as we do not know how the characteristics of the observations in this seven-hour snapshot compared with the same hours on other days, or with other hours on the same day, or with other bilateral rates at the same time.

Table 4.1 D2000-2: Screen at 10:17:40 on 16 June 1993

Currency	Bid	Ask	Quantity Columns	Blank Columns	Latest Price
USD/JPY	106.16	106.25	2 X 1	X X	106.26
DEM/JPY	/	/	//	X X	64.59
USD/CHF	1.4672	1.4679	4 2	X X	1.4676
DEM/CHF	0.8925	0.8933	3 2	X X	0.8929
USD/DEM	1.6439	1.6443	2 1	X X	1.6443
DEM/FRF	3.3633	/	M /	X X	3.3634

Note: USD = U.S. dollar; JPY = Japanese yen; DEM = deutsche mark; CHF = Swiss franc; FRF = French franc.

It is this videotape, initially filmed for its own purposes, that Reuters was kind enough to let us observe, subject to confidentiality commitments. There are four Betacam tapes, which ran virtually continuously, subject to a future minor qualification, from 0832 BST to 1530 BST (on 16 June 1993). The screen does not show the clock time, and the entries are not time-stamped, but a time elapse (time passed since the start of videotaping) was entered onto the tape.¹²

As might be expected, when the commitments made on screen are firm and deals are made at those prices, the original data are, as far as we can judge, remarkably accurate. We ended with only a couple of data points that we felt must be in error. This compares with errors that occur about once in every four hundred entries over FXFX (see Pictet et al. 1994, table 5). By contrast, we are conscious that there will be a number of transcribing errors. In particular, whether because of the need to copy the tapes or for some other reason, the final digit of the five-digit (in one case four-digit) number was often hard to decipher. In particular, it was difficult to distinguish zero from eight when these were faint on the videotape.¹³

In one respect, fortunately, the data are self-checking. When a deal occurs, the transaction price in the right-hand column *has* to be the same as the *prior* (i.e., within seconds earlier) bid, or ask, that was hit and must change the quantity offered at that prior price, and also the price itself, should the quantity be fully taken up. The two series (i.e., of transactions prices, on the one hand, and bid-ask prices and their associated quantities, on the other) were transcribed at

12. We were working at Harvard University when we sought to take the details of the tape, every entry, from the video onto paper and then back onto electronic diskette. Since no Betacam video machines were available in the United States, the tapes were first copied onto S-VHS, and the entries on the S-VHS tapes were viewed over a special video player, with adjustable speeds, forward and backward, pause, etc.

13. The transcription from video to paper was primarily done by the wives of two of the authors, Mrs. Margaret Goodhart and Mrs. Keiko Ito, also with the assistance of Ms. Yoko Miyao, who did this extremely complex and difficult exercise in a dedicated, patient, and conscientious fashion, and we are most grateful to them. But there will inevitably be some errors in variables.

separate times. By marrying these up¹⁴ and reviewing in cases of errors, we can both cross-check the accuracy of our transaction data and get some idea of the remaining errors in variables for the entries (bid-ask and associated quantities offered) where no such cross-check was possible.¹⁵

Turning now to the data themselves, the database divides into two separate parts. First, there is the deutsche mark/dollar market. This is the dominant exchange rate in the foreign exchange market overall, and its dominance of the electronic market in our snapshot is even more marked. There were 799 bid entries and 823 ask entries (note that these entries would usually come from separate banks). Quantities offered at the bid were entered on 802 occasions and at the ask on 841 occasions. (Note that the quantity offered can, and does, change quite frequently without an associated bid-ask price change. Similarly, the price can change without the associated quantity being altered; this happened on more occasions than we would have expected, perhaps because a bank changed the price for a given amount that it wanted to trade.) Although we cannot possibly deduce the total number of *independently* made entries, these might conservatively be put at around fifteen hundred in seven hours, or two hundred or so per hour. This compares with some thirty-five hundred entries over FFX for the deutsche mark/dollar bilateral exchange rate in the same hours, about five hundred per hour. Considering that FFX represents almost costless advertising and is the most commonly used indicative foreign exchange price screen, this shows just how busy the deutsche mark/dollar market on D2000-2 was during this snapshot.

The number of deals in the deutsche mark/dollar was also quite large, relative to the commercial target, reported in section 4.1, of one thousand per day for deals in all fifteen exchange rates. During this snapshot, there were 186 deals done at the bid and 251 at the ask. Whether this ratio of deals to bid-ask entries is high, low, or normal, we cannot tell. We examine whether this ratio varied significantly from half hour to half hour over our data period in section 4.3.

The depth of the deutsche mark/dollar on D2000-2 was fairly good, although it can, and no doubt will, improve further. Following a deal that exhausted the

14. There were a couple of cases when we could not marry the two data points, despite several reviews. It is this to which we referred earlier as the only examples of probable errors in the original data.

15. Thus, the cross-check revealed that the accuracy of visually timing the exact moment of an entry on a screen was to within about plus or minus three seconds. From the adjustments and reviews that had to be made to marry the transaction price data with the bid-ask (and associated quantity) data, it may well be that the final digit in the remaining data is incorrect about once every thirty observations and the penultimate digit incorrect once every one hundred observations. Some of our statistical anomalies, e.g., the few zero and negative spreads and the incorrect direction of price movement following a deal, need to be seen in that context. Such inevitable human error could have been eliminated had the data been available in electronic disk form, but that was not on offer. Moreover, there are some advantages in getting to know the raw data thoroughly before proceeding to econometric testing.

quantity offered or the removal of a bid-ask price, most of the time there was another limit order on the computer at a closely related price. Histograms of quantities offered at the bid measured over both frequency and duration of entry are shown in figures 4.2 and 4.3. The histograms for the ask are nearly identical and have been omitted to save space. From these it can be seen that the frequency and length of time during which *no* bid or ask price is on the screen for the deutsche mark/dollar are both few and brief.

Note that the majority of the quantities offered, both at the bid and at the ask, are usually at or below 5. Consequently, the average size of deal here is also low. We cannot estimate it exactly because we cannot see the actual data lying behind M and L. If, however, we take M to be 8 on average and L to be 15, then the average size of deal at the bid was \$2.51 million and \$2.49 million at the ask, that is, of similar size. A recent paper by Garrett Glass (1994), examining all foreign exchange deals over the Multinet system, puts the aver-




Fig. 4.2 Bid quantity frequency: deutsche mark/dollar




Fig. 4.3 Bid quantity duration: deutsche mark/dollar

age size of deals at about \$9 million.¹⁶ Be that as it may, it is the case that deals in the deutsche mark/dollar D2000-2 market were, by this standard, unrepresentatively small on average. Why this should have been so, we do not know, but Lyons (chap. 5 in this volume) reports that the average size of deals done through brokers is lower than that of customer deals, and his figure for the size of average broker deals is not that much larger than that shown here.

One factor reducing the number/duration of occasions on which there might have been no entry in the deutsche mark/dollar ask series was that a participant, presumably a single bank, kept an off-market ask entry in the computer at 1.6475 when the market was actually running at about 1.6440. When no other entry was better, this was triggered (see fig. 4.4). As the graph shows, the U.S. dollar appreciated sharply thereafter, and the bank involved presumably disposed of its unwanted dollars. In the meantime, however, it represented a nuisance entry for us, distorting the true underlying pattern of the market. No deal was, naturally enough, done at such an off-market price, prior to the occasion of the dollar appreciation. We decided to remove these off-market asks (between observations 250 and 450 on the ask side). We did not remove the few asks at the same price earlier (around the fiftieth observation) since these were not seriously off market (nor did we remove two solitary occasions of off-market *bids* at 1.6405). The resulting, adjusted ask series looks as follows, as shown for comparison in figure 4.5.

As these charts clearly show, the major events in the foreign exchange market on 16 June 1993 were two brief periods of sharp appreciation in the U.S. dollar, the first lasting from about 1339 BST to about 1345 BST and the second from about 1443 BST to 1445 BST, as indicated by the time-stamp on the FXFX data series. The average price of FXFX quote entries in each minute during the course of these two jumps is shown in table 4.2.

The underlying cause, from "news" arrival, of these dollar appreciations against the deutsche mark are clear enough, but their exact timing is difficult to relate to the news items coming over AAMM (the Reuters news page) on that day. The news on that day was "bearish" for Germany and "bullish" for the United States (see table 4.3). Possibly the 1338–1345 BST jump in FXFX could have been triggered by the U.S. housing figures (certainly the dollar opened firm in the United States) and the 1442–1445 BST jump by a delayed reaction to the German government report, but such links cannot be firmly established. The finding here is consistent with other findings in the literature that tend to experience difficulties matching news events to jumps in the asset price, and vice versa. Nevertheless, one can hardly query the time-stamp on the FXFX data, and the extent and timing of these jumps are very closely matched by the data on D2000-2, as will be discussed further below. One inter-

16. Considering that deal size is highly skewed, we wonder whether he meant *median* when he wrote *average* here.




Fig. 4.4 D2000-2 ask data: deutsche mark/dollar




Fig. 4.5 Filtered D2000-2 ask data: deutsche mark/dollar

esting feature of these jumps in the value of the dollar is that they were associated with great activity on the ask side of the market and very little action, even in the guise of price revisions, on the bid. From 1337 BST to 1345 BST, there were seventeen deals at the ask in D2000-2 and none at the bid. Over the same period, there were some thirty-nine price revisions at the ask and thirteen at the bid, two of these remaining established and unchanged for almost two minutes each. From 1442 BST to 1446 BST, there were thirteen deals at the ask and none at the bid. There were some twenty-six price revisions at the ask. A few seconds after the start of the dollar appreciation, the existing bid price was removed from the screen, and for the remaining three and a half minutes of the appreciation no bid price at all was posted: this was the longest gap in having a firm price set for either the bid or the ask in our data set for the deutsche mark/dollar. Otherwise, price setting in the deutsche mark/dollar over D2000-2 was nearly continuous.

A graphic representation of the bid-ask prices quoted, the occasion and price

Table 4.2 Periods of Appreciation of the U.S. dollar

BST	Bid	Ask	Number of Observations
13:38	1.6474	1.6481	7
13:39	1.6486	1.6494	6
13:40	1.6494	1.6503	6
13:41	1.6500	1.6506	5
13:42	1.6503	1.6512	6
13:43	1.6525	1.6535	7
13:44	1.6540	1.6550	6
13:45	1.6553	1.6564	7
13:46	1.6552	1.6560	7
14:42	1.6571	1.6580	8
14:43	1.6575	1.6584	8
14:44	1.6594	1.6602	5
14:45	1.6600	1.6606	5
14:46	1.6601	1.6611	5

Table 4.3 News on U.S. and German Economies

BST	AAMM Report
12:13:18	"German unemployment could top 4 million—Rexrod"
12:34:40	"Next Bundesbank rate cut seen most likely in July"
12:53:12	"German industry says economy still declining"
13:01:44	"German institute sees no recovery before mid-1994"
13:32:04	"US May Housing Starts rose 2.4%"
13:37:04	"US Home Building in May is strongest in 5 months"
13:46:54	"Bonn can live with current mark-dollar rate"
13:56:30	"German Govt source sees no danger for mark"
14:15:40	"Dlr opens firm in US, surges on German comments"
14:20:12	"US May Industrial Output rose, capacity use steady"
14:32:04	"Bonn wants lower short-term rates—Source"
14:33:48	"US May Housing Starts Rise is modest—analysts say"
14:41:08	"Mark falls against dollar after govt comments"

of deals, and the quantities offered for the deutsche mark/dollar, for the first through the seventh hour, is shown in appendix figures 4A.1–4A.22.

Such continuous price setting was not the case for the other five bilateral exchange rates exhibited on the screen during our seven-hour snapshot. Simple observation revealed that market activity in these rates on D2000-2 over our data period was far more patchy. Initially, the rates were not put onto the screen for some sixteen minutes after the deutsche mark/dollar was shown. Thereafter, during the following six and three-quarter hours, there were in some cases quite long gaps in setting bid-ask prices. The average quantities dealt ranged from just over \$1 million (Swiss franc/dollar bid) to nearly \$3 million (French

franc/deutsche mark ask). The data are shown in table 4.4; the figures in parentheses in the table report the original average deutsche mark size when the deals were done in units of deutsche mark 1 million. Deals were, however, much fewer in number than for deutsche mark/dollar. When there are large price movements, the majority of the deals seem to be purchases of the appreciating currency, and the majority of quotes are on the strong side of the market (see table 4.5). We pursue this effect somewhat further in section 4.4 below. Data on these deals and the number of bid-ask price entries are given in table 4.4, and histograms of the bid quantities offered are shown in figures 4.6–4.10; again, the similar ask histograms are omitted to save space.

These histograms show differing patterns. The quantities offered on the dollar-based bilaterals (i.e., deutsche mark/dollar, yen/dollar, Swiss franc/dollar) are predominantly for one or two units, with increasingly few offers made as

Table 4.4 Analysis of Deals and Quotes

	Number of Quotes		Number of Deals		Average Size of Deals*	
	Bid	Ask	Bid	Ask	Bid	Ask
Yen/dollar	93	127	12	17	2.33	1.55
Yen/deutsche mark	99	54	15	2	1.91 (3.15)	2.13 (3.50)
Swiss franc/dollar	142	134	18	33	1.125	1.67
Swiss franc/deutsche mark	121	168	19	45	1.26 (2.08)	2.71 (4.45)
French franc/deutsche mark	98	79	14	11	2.71 (4.45)	2.97 (4.88)

Note: Figures in parentheses report the original average deutsche mark size when the deals were done in units of DM 1 million.

*Based on the assumption that $M = 8$, $L = 15$.

Table 4.5 Relation between Direction of Deals and Currency Change

	Number of Deals		Currency Value		% Change
	Bid	Ask	Start	Finish	
Deutsche mark/dollar	186	251	1.6450	1.6585	+ .82
Yen/dollar	12	17	106.25	106.70	+ .42
Yen/deutsche mark	15	2	64.63	64.27	- .56
Swiss franc/dollar	18	33	1.4690	1.4840	+1.02
Swiss franc/deutsche mark	19	45	.8935	.8953	+ .20
French franc/deutsche mark	14	11	3.3648	3.3623	- .07




Fig. 4.6 Bid quantity frequency: dollar/yen




Fig. 4.7 Bid quantity frequency: deutsche mark/yen




Fig. 4.8 Bid quantity frequency: dollar/Swiss franc




Fig. 4.9 Bid quantity frequency: deutsche mark/French franc




Fig. 4.10 Bid quantity frequency: deutsche mark/Swiss franc

size increases. The quantities offered on the deutsche mark-based bilaterals (i.e., yen/deutsche mark, Swiss franc/deutsche mark, and French franc/deutsche mark) show many more (proportionately) larger offers, quite remarkably so for the French franc/deutsche mark (fig. 4.9). One possible explanation is as follows. Suppose that the European cross-rates tend to move less than the dollar-based bilaterals; then the risk involved in building up inventories for a dealer is less. Hence, a larger unit bid is offered. Now, the Swiss franc/deutsche mark and French franc/deutsche mark rates should move less than the correspondent currencies vis-à-vis the dollar because the deutsche mark and the French franc are in the exchange rate mechanism (ERM), and the Swiss franc closely follows the deutsche mark historically. Even the yen/deutsche mark volatility tends to be less than in the yen/dollar or deutsche mark/dollar rates.

We should again stress that we have no means of knowing whether these, somewhat patchy, results were representative of activity in these exchange rates at other times of the day (note that activity in the yen/dollar exchange rate might be expected to be somewhat muted in European market space) or on

other days or whether they would have been representative of the nine other unshown bilateral exchange rates. Moreover, the use of electronic market systems is developing rapidly over time. Be that as it may, the somewhat occasional nature of the market, then, in these other five exchange rates means that we will concentrate most of our econometric studies on the deutsche mark/dollar.

4.3 Comparison of FAFX and D2000-2

As described in the introduction, indicative screen prices, as provided over FAFX, constitute the basis for almost all current time-series studies of the foreign exchange market. While there is no doubt that these are close enough approximations to the underlying firm quotes for low-frequency studies (e.g., frequencies of one hour or longer), concern has been expressed as to whether they do necessarily provide sufficiently close approximations to the underlying firm data for very high-frequency studies. For example, Baillie and Bollerslev (1991) have conjectured that the negative moving average (MA) characteristics found in FAFX ultra-high-frequency data may be a facet of their indicative nature and that the underlying price(s) would not exhibit this characteristic (see also Zhou 1992; Bollerslev and Domowitz 1993; and Bollerslev and Melvin 1994).

Now that we have a seven-hour snapshot of firm prices in D2000-2, we can, *in principle*, make a comparison of them with the bid-ask series from FAFX over the overlapping period for the three data sets deutsche mark/dollar, yen/dollar, and Swiss franc/dollar. A problem, however, is that the D2000-2 data series is not time-stamped, although it does have a time elapse shown on the videotape. In practice, of course, the two series can be matched pretty closely by eye alone by matching the two occasions of short-term appreciation in the deutsche mark/dollar.

To try to match the series even more closely, we constructed artificial series for both the D2000-2 and the FAFX deutsche mark/dollar, bid and ask, with observations evenly spaced every five seconds. (Note that in both cases the original series is irregularly timed and hence cannot be directly correlated.) We assumed, for the purpose of matching (D2000-2 and FAFX) only, that the existing price held until revised, for the purpose of interpolation, where necessary. When no price was exhibited on D2000-2, we treated the prior price as still holding, *except* for the gap in the bid price in the second jump, discussed in the preceding section, where we applied a linear interpolation (between 1.6565 and 1.6590).¹⁷ Alternative rules of thumb for interpolation could have

17. When we subsequently used this series for econometric work, we changed this rule of thumb so that, when a *deal* exhausted the quantity offered and no price was then shown, we took the *next* reported price as becoming effective. Otherwise, the estimated (absolute) price change, following a deal, would have been biased downward.

been tried, but we are confident that doing so would have made no difference for this timing exercise.

Our crucial assumption is that price adjustments on FXFX and D2000-2 would be synchronous. We believe that to be justified. Studies made by one of us (Goodhart 1989) of the reaction of FXFX bid and ask prices to precisely timed news announcements (e.g., U.S. "news" released at 0830 EST) show that these are virtually instantaneous (a few seconds at most), and we should surely expect no slower reaction where prices represent firm commitments (see, e.g., Ederington and Lee 1993). Accordingly, our strategy was to assume that prices in both series would move synchronously. Given this assumption, our approach was to compare the correlation of the two series for the deutsche mark/dollar as we varied their temporal overlap and see which temporal overlap gave the best fit.

In practice, all the exchange rate action came in the second half of our data period (the last two tapes), and the market was so flat in the opening hours (tapes) that we could not find any clear peak in the fit when starting from the front. We therefore worked from the back, fitting the final tape to the FXFX data, to the front. In the event, and slightly disturbingly, we found a twenty-second discrepancy between our best-fit timing for the comparison of the bid and the ask series (see table 4.6). However, given our exact knowledge of how the bid and ask series are timed relative to each other on D2000-2, we overrode this apparent discrepancy from the time-series fitting exercise and split the difference between the two with the result that the observations on D2000-2 are all properly aligned with each other.

This then gave us the basis for comparison of the D2000-2 bid-ask series with the FXFX series over a closely matched data period (with the exact match uncertain by some fraction of a minute). We have to be careful, however, in using the interpolated five-second series themselves in econometric comparisons since the interpolations distort some of the characteristics of the raw data. There were some eight hundred observations in the basic D2000-2 series and about five thousand in the interpolated series for D2000-2. By construction, the extra forty-two hundred observations will exhibit no change, which must tend to drive any estimated autocorrelation toward zero and may also bias the

Table 4.6 BST: Best Estimated Start Times for Tapes

	Bid	Ask
Tape 4	13:40:47	13:41:07
Tape 3	11:59:10	11:59:30
Tape 2	10:15:37	10:15:57
Tape 1	8:31:40	8:32:00

Note: In each case, the finish of tape $t - 1$ was about one second before the start of tape t . For tape 1, the start time is given from the first quote, of deutsche mark/dollar bid and ask: the tape starts with a blank screen almost exactly eight minutes before.

ARCH characteristics. We discuss some of the issues raised by the question of whether to scale the series by time or by tick activity at greater length in section 4.4 below.

Subject to that condition, the means of the bid-ask in both series (FXFX and D2000-2) are almost identical. A graph of the time path for the deutsche mark/dollar from the two sources looks like one line (see fig. 4.1). Thus, the time path of the indicative quotes *can*, on this evidence, be taken as a very good and close proxy for that in the underlying firm series. Nevertheless, some of the characteristics of the bid-ask series (e.g., the pattern of autocorrelation) are somewhat different. Even so, both series indicate a somewhat similar GARCH pattern. As would be expected, the two series are cointegrated, with the indicative series responding more to deviations from the equilibrium (i.e., a larger and more significant negative coefficient on the error correction mechanism). By contrast, the characteristics of the spreads in the FXFX as compared with the touch in D2000-2 are markedly different. The spreads in the FXFX series show clustering among a small number of standard values (e.g., 5, 7, and 10 pips for the deutsche mark/dollar), whereas the spreads at the touch show no such signs of clustering.

The basic characteristics of the temporally matched, filtered (but not interpolated) series are shown in table 4.7. The main pattern of results shows that the D2000-2 and the FXFX raw series are, in general, remarkably similar for the deutsche mark/dollar.¹⁸

The differences between the first four moments of the various price series (bid, ask, and average of the bid and ask) in either levels or first differences are minor. The FXFX series in levels have a somewhat lower average value (probably owing to a larger proportion of their observations coming in the earlier part of the period; see table 4.7), an insignificantly lower volatility (standard deviation), and marginally higher skewness and kurtosis. The FXFX series in first differences have lower means, by a factor of one and a half in the mean and about two or three in the bid and ask (perhaps again because of more observations when little was happening in the early part of the period). These FXFX differenced series have a lower skewness and a slightly lower kurtosis.

There is, however, a more marked difference in the autocorrelation data. The FXFX series exhibit stronger negative autocorrelation in all cases and at all lags, particularly after the first lag. This is least marked at the first lag of the bid and ask series, where the D2000-2 coefficient is about -0.61 compared with values of -0.62 (bid) and -0.67 (ask) for the FXFX series. In the average series, the first lag value for D2000-2 drops to -0.37 , compared with -0.61 for FXFX. After the first lag, the absolute size of the negative coefficients, and of the t -values, drops much more rapidly for D2000-2 than for FXFX. The first

18. At some future date, we intend to construct similar tables for the raw data for the yen/dollar and Swiss franc/dollar on D2000-2 and FXFX, temporally matched. Time did not allow us to do so at this stage.

Table 4.7

**Statistical Characteristics of the D2000-2 and FXFX Time Series
Compared (deutsche mark/dollar)**

	D2000-2	FXFX
1. Bid, number of observations: ^a	799	3,484
Mean	1.649007	1.6482
SD	.006060	.0058
Skew	.63670	.9392
Kurtosis	-1.31504	2.1507
2. Difference of bid:	798	3,483
Mean	.00000994	.000003646
SD	.000389	.0004012
Skew	.57095	.0845
Kurtosis	9.35931	6.393
Autocorrelation coefficients:		
1	-.6173 (-17.3) ^b	-.6236 (-36.77)
2	-.1437 (-3.44)	-.3488 (-17.49)
3	-.1105 (-2.63)	-.1917 (-9.32)
4	.0031 (.07)	-.0802 (-4.02)
5	.0758 (2.13)	-.0365 (-2.16)
GARCH: ^c		
A ₀	-.000 (-1.48)	-.000 (-.22)
A ₁	-.514 (-16.97)	-.481 (-31.92)
B ₀	.000 (3.89)	.000 (4.10)
B ₁	.198 (3.92)	.116 (6.96)
B ₂	.728 (14.07)	.849 (38.14)
3. Average of bid-ask, number of observations:	1,581 ^d	3,484
Mean	1.649511	1.6486
SD	.006052	.0058
Skew	.55846	.9400
Kurtosis	-1.40521	2.1503
4. Difference of average:	1,580	3,483
Mean	.00000515	.000003646
SD	.000192	.000371
Skew	.45549	.0920
Kurtosis	13.3980	9.1457
Autocorrelation coefficients:		
1	-.366 (-14.52)	-.6094 (-35.91)
2	-.169 (-6.32)	-.3278 (-16.51)
3	-.109 (-4.06)	-.1659 (-8.12)

(continued)

Table 4.7 (continued)

	D2000-2	FXFX
4	-.082 (-3.08)	-.0586 (-2.56)
5	-.043 (-1.72)	-.0045 (-.26)
GARCH: ^e		
A ₀	.000 (3.04)	.000 (1.56)
A ₁	-.179 (-9.19)	-.026 (-1.40)
B ₀	.000 (.23)	.000 (6.75)
B ₁	.536 (38.93)	.268 (9.01)
B ₂	.540 (89.49)	.621 (16.05)
5. Spread, number of observations:	1,556	3,484
Mean	6.8464	7.090
SD	8.0955	2.689
Skew	4.034	2.604
Kurtosis	27.063	39.380
Autocorrelation coefficients:		
1	.4686 (18.44)	-.0118 (-.70)
2	.1098 (3.91)	.0173 (1.02)
3	.1322 (4.72)	.047 (2.81)
4	-.0027 (-.09)	.042 (2.49)
5	.0500 (1.97)	.044 (2.58)
GARCH:		
A ₀	1.4778 (11.18)	.006 (185.8)
A ₁	.6890 (29.68)	.032 (5.58)
B ₀	1.0234 (4.30)	.000 (.70)
B ₁	.6591 (43.67)	.287 (77.40)
B ₂	.6454 (43.84)	.643 (247.02)

Note: *t*-values are given in parentheses.

^aSince the results for the ask series are almost identical to those for the bid, we have omitted the former to save space.

^bThis is the first difference of the level.

^cWe ran the system

$$\Delta x_t = \alpha_0 + \alpha_1 \Delta x_{t-1} + \varepsilon_t, \quad \varepsilon_t | I_{t-1} \sim N(0, h_t),$$

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 \varepsilon_t^2.$$

^dSince the bids and the asks were put in at separate times, the numbers of calculated means and spreads will be approximately equal to the *sum* of the number of bids *plus* the number of asks.

Table 4.8 Error Correction Mechanism

	FXFX Dependent		D2000-2	
	Coefficient	<i>t</i> -Value	Coefficient	<i>t</i> -Value
Lagged:				
-1	-.207	-12.8	-.009	-6.28
Dependent:				
-2	-.184	-11.4	-.004	-2.64
-3	-.136	-8.7	-.002	-1.50
-4	.002	-1.9	-.001	-.93
-5	-.001	-.9	-.007	-4.81
Lagged:				
-1	-.107	-4.25	-.002	-2.25
Independent:				
-2	-.004	-1.95	-.002	-1.92
-3	-.000	-.10	-.002	-1.73
-4	-.003	-1.48	.001	.64
-5	-.004	-1.54	.001	.64
ECM	-.180	-15.47	-.006	-8.44

four lags in FXFX in each case have significant negative coefficients. This is so only for the averaged series of D2000-2, and the sum of the negative coefficients is always considerably greater in absolute size than -1 for FXFX, whereas it is between -0.75 and -0.90 for D2000-2.

We find relatively little difference in the GARCH data, which approximate to IGARCH values, except that the FXFX series for the changes in the average and the level of the spread show less persistence of volatility (a lower B_1 coefficient) than the D2000-2 series.

One of the main findings about the characteristics of the continuous-time foreign exchange indicative quote series was that they appeared to have a negative moving average component. One supposition was that this could be due to the fact that they were indicative, not firm, quotes. Now that we can observe the firm quotes, the negative moving average does appear somewhat attenuated, especially for the average of the bid and ask, but it remains a highly significant feature of the time series.

The main difference between the two series occurs in the case of spreads. The most distinctive difference relates to the numerical pattern of the spread, with the FXFX data showing the spread clustering around certain conventional values,¹⁹ while the D2000-2 spreads, being at the touch with the bid and ask prices being input usually by different banks, show no such clustering. Histograms of the frequency of spreads at various sizes for D2000-2 and FXFX are

19. This has been widely noted (e.g., Bessembinder 1994; and Bollerslev and Melvin 1994) and was more extensively described and analyzed in Goodhart and Curcio (1991).




Fig. 4.11 Deutsche mark/dollar spread frequency: D2000-2 data




Fig. 4.12 Deutsche mark/dollar spread frequency: FXFX data

shown for deutsche mark/dollar in figures 4.11 and 4.12. The yen/dollar and Swiss franc/dollar charts, which show almost identical patterns, are available from the authors.

One feature of the deutsche mark/dollar spreads in D2000-2 (fig. 4.11) is that there are a number of occasions of zero spread; that is, the best bid and the best ask are equal. In FXFX, when the quotes are input by the same bank, a zero spread would signal an input error.²⁰

These comparative tables possibly understate the extent to which the two quote series actually do move together. As shown in figure 4.1, when the two

20. Cohen et al. (1981) have persuasively argued that a dealer should always prefer to transact with certainty at a firm bid (ask) quote rather than set an ask (bid) quote at a zero, or tiny, spread distance from it with no immediate certainty of transaction, so on these grounds a zero spread in D2000-2 may also represent a transcription error or a dealer error; indeed, most of these occasions lasted for only a very few seconds. Nevertheless, we intend to discuss with practitioners whether there may be any rationale for the existence of zero spreads on D2000-2, e.g., asymmetric trading (execution) costs between the two sides, and, until we have done so, we have decided to let these data stand.

Table 4.9 Regressions between FAFX and D2000-2 Series

Left-Hand-/Right-Hand-Side Variables	Constant	Coefficient on:	\bar{R}^2 (SE)	Dickey-Fuller <i>t</i> -Statistic ^a
FAFX mean/2000 mean	-.0018 (.0016)	1.0011 (.0010)	.995 (.0004)	-18.07
2000 mean/FAFX mean	.0101 (.0016)	.9938 (.0010)	.995 (.0004)	-18.04
FAFX bid/2000 bid	-.0267 (.0022)	1.0162 (.0013)	.992 (.0005)	-16.16
2000 bid/FAFX bid	.0397 (.0021)	.9759 (.0013)	.992 (.0005)	-16.16
FAFX ask/2000 ask	.0315 (.0021)	.9810 (.0013)	.991 (.0005)	-17.12
2000 ask/FAFX ask	-.0175 (.0022)	1.0105 (.0013)	.991 (.0005)	-17.11

Note: Standard errors are given in parentheses.

^aMacKinnon critical 1 percent value -3.896.

interpolated series are drawn on the same graph, there appears to be only one line. If we regress the two interpolated series for the deutsche mark/dollar together, after temporal matching, we get the results in table 4.9. As can be seen, the respective series, for the average of the bid-ask and the bids and asks separately, are all strongly cointegrated (as should be expected). Only in one case, however, when the average of the interpolated FAFX series is regressed on the average of the 2000-2 series, do the coefficients take on their ex ante expected values with a constant insignificantly different from zero and the coefficient on the right-hand-side variable insignificantly different from unity. Otherwise, the constants are all significantly different from zero, with the D2000-2 bid on average just above and its ask just below that on the FAFX series. As might be expected, the D2000-2 bid is slightly less variable than its FAFX equivalent, while the D2000-2 ask is a tiny bit more variable (perhaps a reflection of our treatment of outliers in the data?).

Such a finding of strong cointegration enables us, always subject to our prior *assumption* that the two series are synchronous and our temporal matching procedure valid, to examine short-term dynamics and whether a deviation between the two series is corrected primarily by a shift in the FAFX series or in the D2000-2 series. Our hypothesis is that, since the D2000-2 series is the underlying firm series, the indicative FAFX series should adjust to it, rather than vice versa. When, therefore, examining the error correction mechanism (ECM), we expect a large, significant negative coefficient on the ECM when the change in FAFX prices is the left-hand-side variable and a much smaller, possibly insignificant coefficient when the change in D2000-2 prices is the left-hand-side variable. The ECM is taken, as appropriate, from the residuals of the equations in table 4.9.

Taking the average of the bid-asks as our example (the results will not change much for the bid or ask series individually), we ran regressions, as follows:

$$\Delta \text{ average series } 1_t = f(\text{lags } \Delta \text{ average series } 1, \text{lags } \Delta \text{ average series } 2, \text{ECM}).$$

The results can be seen in table 4.8. As expected, both the ECM and the effect of prior changes in the underlying D2000-2 series on the FAFX series are more strongly pronounced than the effect of the FAFX series, or the ECM, on the D2000-2 series, although the latter is still clearly significant, despite being much smaller.²¹

Since time series on transactions (i.e., the number and value of deals) have not been available for the foreign exchange market, variations in either the frequency of entry or the volatility of indicative prices, or some combination of both, have often been taken as a proxy for the volume of unobservable transactions. Here, we examine whether this may have been a good proxy.²² Since we cannot, however, compare the profile of D2000-2 and total market transactions, we will proceed on the presumption that the former *may* be a good proxy for the latter.

For this exercise, we divide our data period into half hours for the deutsche mark/dollar series. We take these periods from the start, with the result that the final period is not quite a complete period. Then we compare both the frequency and the size of deals in each half-hour period (as a percentage of the total) as compared with the frequency of quote entry (as a percentage of the overall number) and relative volatility (the standard deviation of the average of the bid-ask in the subperiod divided by the overall standard deviation). We also examine how the average size of spread related to these variables. The basic results for the D2000-2 and FAFX variables are given in table 4.10. Then simple regressions between these variables were run, as shown in table 4.11.

The results are disappointing for those who would use the indicative FAFX data as a proxy to infer the underlying transactions series. The FAFX volatility series is an excellent predictor of the volatility in the firm quotes of D2000-2 (e.g., [4] in table 4.11); the spread series of FAFX is a mediocre predictor of the spreads on 2000-2, with the latter in this case being on average lower, but *much* more variable, by a factor of nearly five (cf. rows SS and FS in table 4.10, and see eq. [8] in table 4.11). This must raise some doubts about certain

21. Note that the coefficients will, however, be biased downward by the interpolation process, forcing the interpolations to take a no-change value. The *t*-values will be less affected by such time deformation.

22. We cannot, of course, yet observe any time series of *total* market transactions. All we have now is a short snapshot of data on transactions over D2000-2. If the temporal profile of transactions over D2000-2 should be an inaccurate and biased proxy for the total volume of transactions, then the question of whether the indicative FAFX data provide a good predictor of concurrent D2000-2 deals would not have much importance.

Table 4.10 Deutsche Mark/Dollar: Half-Hour Periods

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
S2 deals ^a (SD)	9.8	7.5	4.5	6.7	4.8	5.7	3.4	4.1	5.5	6.2	10.75	11.7	11.7	7.3
S2 deals, average size (SD)	1.42	1.64	1.41	2.14	2.00	1.25	2.31	1.625	2.04	1.83	1.96	1.42	1.57	2.24
S2 frequency ^a price entry (SF)	8.3	8.4	5.4	7.4	6.6	3.7	3.6	4.1	4.7	5.9	11.5	11.2	11.45	7.65
S2 volatility ^b (SV)	.083	.05	.05	.041	.066	.033	.007	.10	.05	.074	.612	.132	.314	.116
S2 spread (SS)	6.41	4.04	5.34	4.32	5.15	4.41	2.51	5.95	4.07	4.82	17.5	8.43	7.17	8.33
FXFX frequency ^a (FF)	8.04	7.68	7.78	7.95	7.02	7.02	5.09	6.42	6.48	7.38	7.32	7.91	7.25	6.65
FXFX volatility ^b (FV)	.110	.070	.061	.057	.062	.042	.079	.091	.059	.058	.696	.152	.374	.140
FXFX spread (FS)	6.65	6.36	6.83	6.60	6.89	7.15	7.32	7.05	7.15	7.08	7.92	7.36	7.80	7.30

^aAs percentage of total.

^bDivided by volatility of whole period.

Table 4.11 Deutsche Mark/Dollar: Half-Hour Relations, D2000-2 and FAFX

Left-Hand-/Right-Hand-Side Variables	Constant	b_1	b_2	b_3	R^2
(1) SF/FF	-6.8 (-1.1)	1.95 (2.32)			.25
(2) SD/FF	-6.9 (-1.1)	1.96 (2.32)			.25
(3) SD/SF	.4 (.8)	.94 (9.12)			.86
(4) SV/FV	-0 (-.7)	.88 (27.11)			.98
(5) SD/SV	5.7 (7.14)	11.30 (3.95)			.36
(6) SD/SV, SF	.3 (.3)	-.46 (-.17)	.96 (6.41)		.85
(7) SD/FF, FV	-6.9 (-.15)	1.78 (2.83)	9.30 (3.33)		.59
(8) SS/FS	-31.1 (-2.3)	5.27 (2.80)			.34
(9) SS/SV	3.7 (7.7)	21.23 (8.67)			.85
(10) FS/FV	6.8 (64.4)	1.79 (3.82)			.51
(11) SD/FF, FV, FS	-33.0 (-1.9)	2.45 (3.34)	3.48 (.75)	3.11 (1.54)	.64
(12) SD/SF, SV, SS	.3 (.2)	.96 (6.06)	-.64 (-.11)	.01 (.03)	.84

Note: *t*-statistics are given in parentheses. Initial F stands for FAFX series; initial S for System D2000-2. Second letter F represents frequency of quote entry; D is number of deals; V is volatility; and S is spread. So SF is frequency of quote entry over System D2000-2; FF is frequency of quote entry over FAFX; SD is the number of deals on D2000-2; SV is the volatility of D2000-2, etc.

aspects of the results of recent empirical studies based on FAFX data (e.g., Bollerslev and Melvin 1994; and Bessembinder 1994). This is discussed further in section 4.4.2. The frequency of *quotes* series on FAFX was a relatively poor predictor of the quote frequency on D2000-2. Unfortunately, the importance of these series as a predictor of deals is largely in reverse order in this data set. As can be seen (eqq. [3], [6], and [12]), the frequency of quote entries over D2000-2 is the dominant predictor of the number of deal entries, with neither volatility (whose coefficient was even wrong signed) nor spreads being significant. But FAFX entry frequency is a poor predictor of D2000-2 quote entry frequency. Thus, using FAFX data to predict the number of D2000-2 deals was *not* very successful. The frequency of entry (FAFX) was the most significant variable for predicting D2000-2 deals of the data series available over FAFX (eqq. [2], [7], and [11]), but both FAFX volatility and spreads made some positive contribution. We are fully aware of the small size of this sample among many dimensions, length of time, number of observations, etc.

While more work is undoubtedly needed, we must warn that this preliminary exercise suggests that it would be dubious to try to infer transaction frequency from the more widely available FAFX indicative quote data.²³

To sum up, in this section we have sought to compare the characteristics of the D2000-2 and FAFX series over a temporally matched period. The main result is that the time paths for the prices quoted over the two series are *extremely* close and that most of the time-series characteristics of the two quote series are closely similar. The negative autocorrelation is somewhat attenuated, but still highly significant, in the firm D2000-2 series. As expected, the distribution of spreads is markedly different between the indicative series, which clusters at certain round numbers, and the touch with a much more even distribution.

The size of spreads and the frequency of quote entry showed much more variation between subperiods in the D2000-2 series than in the FAFX, and the latter were not good predictors of their D2000-2 counterparts, unlike FAFX volatility, which like its mean value matched D2000-2 almost exactly. This meant that the FAFX data proved to be poor predictors of the frequency of deals over D2000-2 for the deutsche mark/dollar since this was most closely associated with the frequency of quote entries in that same data set.

4.4 The Interaction of Transactions and Bid-Ask Quotes on the Foreign Exchange Market

4.4.1 Characteristics of Transactions Data

In the preceding section, we asked how accurate a proxy the commonly available FAFX data were to the underlying firm D2000-2 quotes (excellent as a guide to price movements) and to the spreads and number of underlying transactions over the same data set (which suggested that a lot of caution would be needed). In this section, we test certain hypotheses about the determinants of the occurrence and size of such transactions and their effect in turn on quote revision. We concentrate solely on the deutsche mark/dollar series here because only in this series are there sufficient data points.

Our first hypothesis is that the time series for transactions prices (returns) will be random walk. This is the standard efficient markets hypothesis. Most of the evidence of autocorrelation in returns in stock markets has related to discrete break points in markets, that is, market openings and closings, weekend effects, and end-tax-year effects (see, e.g., Dimson 1988; McInish and Wood 1991; Wood, McInish, and Ord 1985; Griffiths and White 1993). The

23. We also ran a similar exercise, using hourly data, for the Swiss franc/dollar series, but, with only fifty-one deals in our data period, this was too affected by small sample problems to provide a useful cross-check. Data on this are available from the authors.

foreign exchange market exhibits fewer discrete break points; in any case, our sample is far too small, covering no such break points, to hope to test for any such anomalies.

We exhibit the characteristics of the transactions data separately for transactions at the bid and the ask and also for the two series taken together (to see what the effect on the characteristics would be if, counterfactually, we could not distinguish between deals at the bid and the ask; see table 4.12). During our short snapshot, the deutsche mark/dollar traded upward (i.e., the dollar

Table 4.12 Transactions in Deutsche Mark/Dollar

	Bid	Ask	Bid + Ask
Number	186	251	437
Average size	\$2.51 mn	\$2.49 mn	\$2.5 mn
Levels mean	1.64946	1.64978	1.6496
SD	.0062	.0061	.0061
Skew	.5541	.5571	.5518
Kurtosis	-1.4346	-1.3617	-1.3910
First Differences			
Mean	.000042	.000034	.000019
SD	.00054	.00030	.000269
Skew	5.096	-0.575	1.273
Kurtosis	38.556	10.164	15.326
Autocorrelation coefficient:			
-1	-.084 (-1.11)	-.086 (-1.32)	-.1406 (-2.90)
-2	-.069 (-.90)	.138 (2.13)	.0949 (1.96)
-3	.155 (2.07)	-.085 (-1.30)	.0185 (.38)
-4	-.009 (-.12)	.050 (.77)	.054 (1.12)
-5	.003 (.04)	.042 (.66)	.026 (.54)
GARCH:			
A ₀	-.000 (-.63)	.000 (1.70)	.000 (.52)
A ₁	-.004 (-.55)	-.159 (-1.89)	-.365 (-17.49)
B ₀	.000 (2.89)	.000 (4.11)	.000 (0.32)
B ₁	.415 (20.26)	.572 (4.21)	.553 (19.49)
B ₂	.491 (44.87)	.246 (2.47)	.478 (56.46)
Dickey-Fuller test with 5 lags	-643.34	-91.22	-301.3

Note: *t*-statistics are given in parentheses.

appreciated). So the mean change on all three series was positive, but less so for the composite series because of bouncing between deals done at the bid and the ask.

Because of that same bounce, the absolute size of the negative autocorrelation on the first lag becomes larger (almost doubles) and becomes significant. Thus, we claim to be able to document here the statistical effect of the bounce. It would be possible to use these data to check the accuracy of the Roll (1984) model whereby the size of the bid-ask spread is estimated using only transaction prices. We leave that for later work, although we doubt whether that model would perform well, for example, because the direction of deals is autocorrelated and information asymmetry (volatility) is time varying. The positive coefficients at higher lags on the other two series *may* be owing to the large jumps in the dollar during our short data period. Bollerslev and Domowitz (1993, 1430–32; see also Bollerslev and Domowitz 1991) generate artificial transactions series from automated trade execution algorithms that exhibit positive first-order serial correlation; we find no sign of that outcome in our data set of actual transactions prices.

Hasbrouck and Ho (1987) find that, for the NYSE, “the pattern consists of a large negative auto-correlation at the first lag, followed by positive auto-correlations of decreasing magnitude that are statistically significant . . . through the fifth lag. The negative first order auto-correlation in transactions data is consistent with the findings of other studies. The positive auto-correlations, however, are (in transactions data) new” (1039). While the size and significance of our coefficients are considerably less, the general pattern in our data is exactly the same. With the significant negative first-order auto-correlation being caused by the bounce and none of the later positive autocorrelations being either large or significant, our results are, not surprisingly, consistent with efficiency.

The Dickey-Fuller test indicates stationarity. This does not disturb us. The random walk characteristic of asset prices results from their subjection to a sequence of “news” shocks. At any one point of time, the market price of an asset should have an equilibrium value, dependent on assessments of past “news” shocks. If the time period is short enough, here only seven hours, the amount of additional “news” is limited, so, over very short time periods, one might expect to observe stationarity.

What we do feel remains to be clarified and modeled is the nature of the interaction between a quotes series that shows clear evidence of a negative moving average component and a transactions series that exhibits no such significant autocorrelations. This is the subject of our ongoing research.

According to the simplified models wherein a single dealer undertakes one transaction of a standardized size per period, the dealer should adjust prices until the expectation of a transaction at the bid next period is equal to one at the ask. So the sequence of deals between bids and asks should be random (see, e.g., Admati and Pfleiderer 1988, 1989; and Hasbrouck and Ho 1987). If

inventory effects are present, the sequence might be expected to show some negative autocorrelation. With many dealers posting limit orders and multiple orders possible in any finite period, we would, however, not expect that. Instead, we would expect runs of deals of each kind. We test that hypothesis, both by a histogram showing the lengths of sequences of deals of both kinds and by a formal runs test.

The histogram, figure 4.13, shows that there are a number of runs of deals, at both the bid and the ask, that are *much* longer than one might normally expect to see. These are shown in table 4.13, together with their individual expected probability of occurrence. The probability of finding all such runs together is infinitesimal. Thus, rather like the kurtotic characteristic of the price change series, the run series for deals appears to have a fat tail. As noted earlier, there are indications that runs of similarly signed deals occur when the price series is trending in one direction, for example, dollar buying at the ask where the dollar is appreciating. We show the associated change in the relevant quoted price during each run over the same period in table 4.13.

The formal runs test that we use is the Geary test. This concentrates attention on whether the number of runs observed in the sample is large or small relative to the number that one would expect to occur in a strictly random sample. According to this test, we are led to reject strongly the null that successive observations are independent since the test statistic is -7.11 compared to the standard normal critical value of -2.58 under the null.

Some earlier empirical work has also found evidence that deals tend to run in sequences (bid deals followed by bid deals and ask deals followed by ask deals), for example, Hasbrouck and Ho (1987) and Lease, Masulis, and Page (1991) for the NYSE. Some of the reasons for this are straightforward, for example, a trader with a large order working up the limit order book. We would, however, conjecture, but have yet to do the work required to demon-




Fig. 4.13 Deutsche mark/dollar deal runs: bid and ask combination

Table 4.13 Deal Runs, Price Changes, and Sample Probability of Occurrence

Run Length	Side of Market	Percentage Price Change	Sample Probability
21	Ask	.272	8.8×10^{-6}
17	Ask	.516	8×10^{-5}
15	Ask	.103	2.4×10^{-4}
14	Ask	.122	4.2×10^{-4}
12	Bid	-.030	3.6×10^{-5}
11	Bid	-.097	8.4×10^{-5}
9	Ask	.121	6.8×10^{-3}
8	Ask	.055	.0018
8	Ask	.018	.0018
7	Bid	-.090	.0025
7	Bid	-.055	.0025
7	Ask	.018	.0205

Note: Percentage price change represents the percentage difference in the quotes on the stated side of the market at which the first and last transaction in each run took place. Sample probability is simply the probability, given the sample frequency of each type of deal, of observing n successive transactions on one side of the market, assuming that they are independent events.

strate, that the extent of autocorrelation revealed here is considerably beyond the explicable on the basis of such simple microstructural factors.

The only theoretical explanation yet given for such positive autocorrelation is by Admati and Pfleiderer (1989). They suggest that market dealers may shade the costs of dealing, on one side of the market, to encourage liquidity traders to bunch together on that side, isolating and identifying informed traders on the other: "The intuition behind our results suggests that there will be periods in which prices rise at a slow rate when shares are purchased but fall at a more rapid rate when shares are sold. These periods will be periods of concentrated buying—periods in which it is expected that discretionary *buyers* will be trading" (Admati and Pfleiderer 1989, 209). Our results are very different. In our data, buys concentrate together when prices are rising rapidly and spreads rising, but not enough to choke off the stream of purchases. At such moments, seller-initiated trades dry up altogether. Further research to check whether our results are typical of the foreign exchange market and, if so, what the reasons for this might be would be desirable.

4.4.2 The Interrelations between the Data Series

Given the existence of such long runs of deals at the bid and ask, one variable that may help predict the occurrence of a deal at the bid (ask) is whether there has been a prior deal at the bid (ask). Hence, we now turn to regression analysis to explore the interrelations between our series, separately for both D2000-2 and FXFX. For this purpose, we used our constructed five-second data set, where for D2000-2 a nonentry at either the bid or the ask is replaced by the prior entry, if no deal had occurred, or the subsequent entry following a

deal. There was never more than one deal in any five-second period, but, of course, over longer periods (e.g., one minute) there were often several deals.

For D2000-2 we had the data series shown in table 4.14 for both the bid and the ask; bid series are given the notation B and ask series A. There were thus seventeen basic series for D2000-2, eight bid, eight ask, and the spread. Initially, we used our five-second database, with lags covering the previous two thirty-second intervals and the two minutes before then, for example,

$$BD_{1-6}, BD_{7-12}, BD_{13-24}, BD_{25-36},$$

noted as $BD_6, BD_{12}, BD_{24}, BD_{36}$. In some cases, for example, for spreads and quote revisions, we also used shorter-unit (five-second) lags, noted as Lag 1, Lag 2, Lag 3, etc.

For FXFX, we did not have the first three series, (BD, BDQ, BDE) or QB, so there were four basic series in this case, with similar notation (DB, BF, ADB, and BV), for bid quotes, four for asks, and the spread. This meant that we had over eighty-five basic series (including lags) for D2000-2 and a data set of five thousand observations.²⁴ Our basic approach was to regress each variable of interest on lagged values of all the variables (including the lagged dependent) separately and then include significant values from these first-stage equations in a larger equation to search for the best-fitting equation.

There is a general problem in such exercises of how to scale the data. The two main alternatives are to use standard clock time or transactions (tick) time, whereby each activity observation is ordered consecutively, irrespective of the varying time gap between them. With very high-frequency series, for example, five-second intervals as here, a problem with the use of clock time is that most observations of price changes, deals, etc., are zero. Hence, the distribution of these variables is nonnormal, with a spike at zero. On the other hand, there are certain questions relating to the temporal relations between series, especially in multivariate analysis, that can be answered only using a clock-time scale. Several analysts have wrestled with this problem, notably McNish and Wood (1990, esp. sec. 4.4) with respect to the NYSE and the various studies undertaken by analysts at Olsen and Associates (e.g., Müller et al. 1990; and Dacorogna et al. 1993) of the foreign exchange market. Most empirical work in both stock and foreign exchange markets, has been performed on an activity scale, utilizing tick-by-tick data. The studies (e.g., on price scaling laws), notably those carried out by Olsen and Associates, do suggest that this is probably preferable, where feasible, for the question under consideration. In our case, however, we are interested in multivariate intertemporal relations, so we have primarily used a clock-time scale but have, in certain cases, checked the result from these exercises against similar exercises on an activity scale.

24. Our computer could not handle a general to specific exercise with parameters of this size, although there was relatively little multicollinearity or autocorrelation (apart from the spread, S, which was strongly positively autocorrelated in D2000-2). We ran a simple cross-correlation matrix, which is too large to reproduce but is available from the authors.

Table 4.14 Codes for Variables

	Bid	Ask
Number of deals in period	BD	AD
Quantity traded in deal	BDQ	ADQ
Dummy if deal exhausted quantity	BDE	ADE
Change in quote	DB	DA
Quantity quoted	QB	QA
Frequency of quote revision over period	BF	AF
Absolute value of change in quote	ADB	ADA
Standard deviation of changes in quotes	BV	AV
Spread		S

The following exercises are quite detailed. The relevant tables are tables 4.15–4.32 below. Readers may prefer to skip first to the figures showing, qualitatively, the main directions of relations (figs. 4.14a–c below) and also to the summary of main findings in section 4.4.5 before deciding how much detail in the next few pages they want to absorb.

There were only some 186 bid deals in the deutsche mark/dollar during the more than five thousand five-second intervals. So, to examine the likelihood of a bid deal occurring, we used probit analysis. Our “best” equations for the probability of bid and ask deals occurring are shown in table 4.15.

The main finding from this, which was foreshadowed in the results in table 4.11 above, is that the most important set of variables to determine bid (ask) deals is the *frequency* of bid (ask) quote revisions in the previous few minutes. This frequency, we believe, is probably a proxy for the extent of prior information. When lagged values of BF (AF), the frequency variable, are entered, lags of the dependent variable BD (AD) lose most of the significance they had when entered alone. Besides this frequency variable, in both cases, if there was a deal of the *opposite* sign (e.g., AD6 in the BD equation) in the previous thirty seconds, there is *less* likelihood of seeing a deal now. Bid deals in the deutsche mark/dollar are considerably more likely to occur where current spreads are low (i.e., prices are good) and when prices have recently been improving (DB6 is positive). This suggests that traders are doing their job effectively (i.e., hitting comparatively good prices). A comparison of average spreads when there is no deal and when there is a deal for the deutsche mark/dollar and the yen/dollar is shown in table 4.16.²⁵

The AD (ask deal) results are more problematic, with some nonintuitive variables entering significantly, that is, a *positive* lagged spread (thirty seconds previous), positive changes in *bid* quotes, and a deal quantity variable, ADQ. We surmised that these results might have been due to many of the ask deals

25. Note that the split of the period into subdivisions differs slightly between table 4.10 and table 4.16.

Table 4.15 Probability of Observing Deals

A. Bid				
Probit estimates:				
Number of observations = 4,980				
$\chi^2(7) = 96.1$				
Prob. > $\chi^2 = .000$				
Pseudo $R^2 = .061$				
Log likelihood = -732.90769				
Bid Deal	Coefficient	SE	<i>t</i>	<i>P</i> > <i>t</i>
bf6	.1663953	.0329075	5.056	.000
bf12	.1275713	.0328465	3.884	.000
bd36	.1151307	.0427033	2.696	.007
db6	154.4183	81.54656	1.894	.058
s	-431.2308	91.80828	-4.697	.000
ad6	-.2524728	.1295843	-1.948	.051
_cons	-1.893544	.069593	-27.209	.000
B. Ask				
Probit estimates:				
Number of observations = 4,980				
$\chi^2(10) = 96.2$				
Prob. > $\chi^2 = .000$				
Pseudo $R^2 = .048$				
Log likelihood = -952.29651				
Ask Deal	Coefficient	SE	<i>t</i>	<i>P</i> > <i>t</i>
af6	.0737124	.0311327	2.368	.018
af12	.0777921	.030167	2.579	.010
adq6	.0508467	.0205937	2.469	.014
Lag6 s	88.19644	39.06158	2.258	.024
bd6	-.212043	.0811112	-2.614	.009
bd24	.1275771	.0381741	3.342	.000
db6	107.5494	68.76905	1.564	.118
db12	197.6548	66.76211	2.961	.003
db24	154.0204	58.16457	2.648	.008
db36	116.271	54.50788	2.133	.033
_cons	-1.940738	.0564778	-34.363	.000

Note: bd = bid-side deals; ad = ask-side deals.

occurring in the latter part of the period, when spreads and volatility were high and both bid and ask quotes prices rising markedly. In order to test this, we divided our sample into two parts, the flat first half (observations 1–3560) and the upward-trended second half (observations 3561–5000), and redid the probit analyses for both the bids and the asks. The results for bid deals remained much the same. For ask deals, the spread becomes negative (as expected) in

Table 4.16 Spreads: A Comparison of Spreads at Ordinary Times with Those at Transaction Times

A. Deutsche Mark/Dollar Bid-Ask Spread						
Hour	Bid-Ask, All Samples			Bid-Ask, Transaction Time Only		
	Mean Unit = DM/\$	Median Unit = DM/\$	Number of Observations	Mean Unit = DM/\$	Median Unit = DM/\$	Number of Observations
0	.0004214	.00030	607	.0004125	.00020	72
1	.0004992	.00040	708	.0003587	.00030	46
2	.0003791	.00030	671	.0002391	.00020	46
3	.0005388	.00040	577	.0003700	.00030	30
4	.0005110	.00040	647	.0003113	.00030	44
5	.0011005	.00070	656	.0010630	.00060	92
6	.0007651	.00070	602	.0004777	.00040	72
7	.0007530	.00050	49	.0004000	.00040	6

B. Yen/Dollar Bid-Ask Spread						
Hour	Bid-Ask, All Samples			Bid-Ask, Transaction Time Only		
	Mean Unit = Yen/\$	Median Unit = Yen/\$	Number of Observations	Mean Unit = Yen/\$	Median Unit = Yen/\$	Number of Observations
0	.11213	.15000	136	.01000	.02000	2
1	.10996	.13000	720	.08000	.09000	3
2	.14967	.20000	720	.08833	.09500	6
3	.18212	.20000	720	.14000	.14000	2
4	.15825	.19000	554	.04500	.05000	4
5	.14814	.15000	199			0
6	.10598	.10000	112	.08333	.03000	6
7	.08000	.08000	14			0

Note: (1) Each hour has a maximum of 720 observations (five-second intervals). If an ask or bid is missing, then that bracket is not counted in the left-hand-side panels of "all" observations. (2) Transaction time bid-ask spread is the bid-ask spread of the five-second bracket, preceding the five-second bracket where a transaction occurs. There are instances where transactions occur even without one of the bid or ask being shown on the screen (just before the transaction is recorded). These are treated as missing observations in the right-hand-side panels.

the first half and insignificant in the second part; and the change in the ask price (DA36) also enters negatively, as expected, in the first half of the period. Apart from the insignificant spread, the results for the first part ask are similar to those of the contemporaneous bid. The table giving these two half-period results is available on request from the authors.

Overall, however, the fit was rather poor. Perhaps it was expecting too much of the data to be able to predict the probability of a deal within a period as short as five seconds. So we lowered the frequency of analyzed periodicity to a minute. Within a minute, however, there were often several deals. So we used ordered probit analysis to estimate the interrelations. Somewhat to our surprise, the change of periodicities to the lower frequency of one-minute inter-

vals made relatively little difference to the major apparent patterns of relations (see table 4.17).

Given the probability of a deal, the next question is what will be the volume, the size of the deal. In 145 of 186 deals at the bid and 179 of 251 deals at the ask, the deal, however, exhausted the outstanding quantity offered. So the size of the deal was usually limited by the amount on offer. That means that it is more sensible to try to model the amounts offered by the dealers (BQ and AQ) than the amounts sought by the hitters (i.e., the supply function is better identified than the demand function).

Similarly, of course, the price of the deal has to be at the price posted, either the bid or the ask, in the firm quotes. So we turn next to an analysis of the determinants of the changes in such prices, DA and DB. As noted earlier, when a quote is hit and exhausted, the price *must* change to the next limit order, if such exists. There is also known to be negative autocorrelation in the quote series. Our first basic exercise was, therefore, to regress DA and DB against their first six, $t-1$ to $t-6$, own lags and the dummy exhaust variable, BDE and ADE, taking the value 1 when the quote was exhausted by a deal. The results

Table 4.17 Ordered Probit Analysis on Data at One-Minute Intervals

A. Bid-Side Deals				
Number of observations = 403				
$\chi^2(4) = 38.1$				
Prob. $> \chi^2 = .000$				
Pseudo $R^2 = .056$				
Log likelihood = -320.24				
Bid-Side Deals	Coefficient	SE	t	$P > t $
bf6	.0357	.0133	2.67	.008
bf24	.0259	.0079	3.28	.001
db24	64.45	31.34	2.05	.041
Lag1s	-180.0	89.27	-2.02	.044
B. Ask-Side Deals				
Number of observations = 391				
$\chi^2(4) = 34.4$				
Prob. $> \chi^2 = .000$				
Pseudo $R^2 = .040$				
Log likelihood = -408.15				
Ask-Side Deals	Coefficient	SE	t	$P > t $
af6	.0351	.0117	2.993	.003
ada6	198.20	59.49	3.331	.000
da12	-69.07	49.96	-1.383	.168

Note: db = change in bid quote; da = change in ask quote.

are as shown in table 4.18. The value of the dummy exhaust variables (BDE and ADE) was in each case about $\pm .000375$, showing that this is the average price revision (down following a bid exhaust, up after an ask exhaust), or alternatively the gap between limit orders, following a deal. The negative values for the lagged own values are consonant with the now-well-established high-frequency negative autocorrelation.

The lower value of the coefficient on the first lag than in table 4.7 above is due to the fact that the series here are on clock time, five-second intervals, and not taken, as in table 4.7, by consecutive quotes. Consequently, most of the observations on price changes show zero. When we reran the exercise on ex-

Table 4.18 Basic Determinants of Quote Revision

A. Bid Quote Revision				
Number of observations = 4,976				
$F(7,4969) = 87.9$				
Prob. > $F = .000$				
$R^2 = .110$				
Adjusted $R^2 = .109$				
Root MSE = .0002				
Change in Bjd Quote (db)	Coefficient	SE	t	$P > t $
Lag1 db	-.11656	.0136231	-8.556	.000
Lag2 db	-.1176993	.0137196	-8.579	.000
Lag3 db	-.1320021	.0138006	-9.565	.000
Lag4 db	-.0546471	.0137991	-3.960	.000
Lag5 db	-.0210431	.0137147	-1.534	.125
Lag6 db	-.0776679	.0136245	-5.701	.000
Lag1 bde	-.0003729	.0000189	-19.745	.000
_cons	.0000157	3.21e-06	4.881	.000
B. Ask Quote Revision				
Number of observations = 4,976				
$F(7,4969) = 114.1$				
Prob. > $F = .000$				
$R^2 = .138$				
Adjusted $R^2 = .137$				
Root MSE = .0002				
Change in Ask Quote (da)	Coefficient	SE	t	$P > t $
Lag1 da	-.111342	.0134768	-8.262	.000
Lag2 da	-.08533	.0133953	-6.370	.000
Lag3 da	-.0669398	.0134363	-4.982	.000
Lag4 da	-.0322129	.0134368	-2.397	.017
Lag5 da	-.1609648	.0133957	-12.016	.000
Lag6 da	-.0400806	.0134622	-2.977	.003
Lag1 ade	.0003769	.0000161	23.339	.000
_cons	-8.79e-06	3.06e-06	-2.871	.004

actly the same basis but omitting those observations when prices changes were zero, we got the results shown in table 4.19. The absolute size of the coefficients of the lagged dependent variables increases by a factor of about five times (as the 80 percent of zero observations in the complete, clock-time, sample are removed), but the standard errors increase by as much, or slightly more, so the t -values actually decline, just, on balance. Since there virtually has to be a change in price after a deal exhausts the previous quote entry, coefficients of the deal exhaust dummies, BDE and ADE, rise only slightly, and, with a commensurately higher standard error, their t -values fall from around 20 to

Table 4.19 Basic Determinants of Quote Revision: Zero Changes Omitted: Tick by Tick

A. Bid Quote Revision				
Number of observations = 727				
$F(7,720) = 34.2$				
Prob. > $F = .000$				
$R^2 = .249$				
Adjusted $R^2 = .242$				
Root MSE = .0005				
Change in Bid Quote (db)	Coefficient	SE	t	$P > t $
Lag1 db	-.4204473	.0687554	-6.115	.000
Lag2 db	-.4810583	.0717263	-6.707	.000
Lag3 db	-.5284103	.0692662	-7.629	.000
Lag4 db	-.2811608	.0843628	-3.333	.000
Lag5 db	-.0902551	.0788967	-1.144	.253
Lag6 db	-.4268289	.0869281	-4.910	.000
Lag1 bde	.0004502	.0000513	-8.783	.000
_cons	.0000989	.0000223	4.432	.000
B. Ask Quote Revision				
Number of observations = 747				
$F(7,740) = 43.1$				
Prob. > $F = .000$				
$R^2 = .289$				
Adjusted $R^2 = .283$				
Root MSE = .0004				
Change in Ask Quote (da)	Coefficient	SE	t	$P > t $
Lag1 da	-.7198426	.0813796	-8.845	.000
Lag2 da	-.3373846	.0585086	-5.766	.000
Lag3 da	-.4080851	.0765337	-5.332	.000
Lag4 da	-.2991646	.075899	-3.942	.000
Lag5 da	-.53916	.0578206	-9.325	.000
Lag6 da	-.1875199	.0644919	-2.908	.004
Lag1 ade	.0004132	.0000445	9.295	.000
_cons	-.0000571	.0000209	-2.729	.007

about 9. The resultant series without the zeroes (i.e., in transaction time) is much more variable, so, although the fit of the series is much improved (the adjusted R^2 doubles from around .12 to about .25), the root MSE also doubles.

We then explored to find other variables that might contribute significantly to the determination of quote revision, although the own lags out to $t-6$ and the exhaust dummy remained the key variables. The main additional variables that entered in table 4.20 were the *spread* with a one-period lag, negatively for the ask and positively for the bid (i.e., where the spread was unusually large, someone would come forward with a more competitive quote); longer own lags (although this was more apparent in equations run *without the spread*, as shown in table 4.21); and some *volatility* variables.²⁶

When the spread variable is *not* included, changes in the ask price have a strong positive effect on changes in the bid price, whereas changes in the bid price had a weaker effect on changes in the ask prices (see the coefficients italicized in table 4.21). But the sum of the coefficients is *well below* unity. What this means is that, in this market, a change in the best bid (ask) has only a slight effect on the contemporaneous ask (bid). Most of the immediate effect becomes translated into a changed spread, which is highly positively autocorrelated. The spread returns toward normal only slowly. So, in this market, with best bids and asks being entered by different banks, the hypothesis that these two quotes will be revised closely and quickly in step with each other is convincingly refuted; instead, bids and asks vary somewhat independently, rather like two variables that are cointegrated in the longer run, with the spread acting as the error correction mechanism between them.

We have no convincing explanation for the asymmetry whereby the change in the ask quote price had a stronger effect on the bid quote price than vice versa. We initially thought that this might be due to the surge in the value of the dollar in the second half of the period, affecting first ask deals and quotes and thereafter bid quotes, but, when we divided the period into two and reran, this hypothesis was refuted since, although the effect of DA on DB was slightly weaker than in the full sample, it was clearly stronger in the first, untrended part of the period than in the second part, when the dollar strengthened.

We also looked for any signs that either the event or the size of deals influenced quotes, apart from the exhaust dummies, which, as already noted, were highly significant. We found generally rather weak effects, as in table 4.20 below, of these variables on quote revisions, but where significant usually of the expected sign. Thus, in some of the equations for bid quote revisions, DB, the event (BD) or the quantity (BDQ) of a deal in prior periods would enter

26. Such volatility variables were usually AV12, or sometimes BV60, in the ask price change, DA, equation and ADB in the bid price change, DB, equation. Rather counterintuitively, this latter variable was positive in the DB equation, and, when it entered, AV12 was negative in the DA equation, implying that higher volatility led to finer, more competitive prices being posted, but the significance level of this is not high.

Table 4.20 The Determinants of Quote Revision

A. Bid Quote Revision				
Number of observations = 4,976				
$F(13,4963) = 64.4$				
Prob. > $F = .000$				
$R^2 = .144$				
Adjusted $R^2 = .142$				
Root MSE = .0002				
Change in Bid Quote (db)	Coefficient	SE	t	$P > t $
Lag1 db	-.1199148	.0139002	-8.627	.000
Lag2 db	-.127044	.0139954	-9.078	.000
Lag3 db	-.1449427	.0140143	-10.342	.000
Lag4 db	-.0792908	.0141885	-5.588	.000
Lag5 db	-.0478669	.0140588	-3.405	.253
Lag6 db	-.1040509	.0140001	-7.432	.000
Lag1 bde	-.0003474	.0000196	-17.734	.000
db12	-.0503687	.0076551	-6.580	.000
bdq6	-.0000108	3.32e-06	-3.240	.001
adb6	.0375697	.0086841	4.326	.000
adb24	.0103388	.00735	1.407	.160
adq24	5.36e-06	1.70e-06	3.148	.002
Lag1 s	.0375859	.0048246	7.790	.000
_cons	-.0000196	4.76e-06	-4.116	.000
B. Ask Quote Revision				
Number of observations = 4,975				
$F(11,4964) = 76.5$				
Prob. > $F = .000$				
$R^2 = .145$				
Adjusted $R^2 = .143$				
Root MSE = .0002				
Change in Ask Quote (da)	Coefficient	SE	t	$P > t $
Lag1 da	-.0946171	.0140177	-6.750	.000
Lag2 da	-.0699815	.0138121	-5.067	.000
Lag3 da	-.053245	.0138313	-3.850	.000
Lag4 da	-.018927	.0138015	-1.371	.170
Lag5 da	-.1497416	.0136565	-10.965	.000
Lag6 da	-.0278212	.0135895	-2.047	.041
Lag1 ade	.0003646	.0000176	20.663	.000
ad6	.00001	5.83e-06	1.723	.085
16db	-.0307674	.0127157	-2.420	.016
bv60	.0000389	.0000135	2.881	.004
Lag1 s	-.0278851	.0048746	-5.720	.000
_cons	-2.63e-06	4.68e-06	-.563	.574

Table 4.21 Quote Revisions

A. Bid Reaction to Changes in Ask Quotes

Number of observations = 4,975
 $F(16,4959) = 48.1$
 Prob. > $F = .000$
 $R^2 = .134$
 Adjusted $R^2 = .131$
 Root MSE = .0002

Change in Bid Quote (db)	Coefficient	SE	<i>t</i>	$P > t $
Lag1 db	-.1428408	.0136389	-10.473	.000
Lag2 db	-.1473166	.0137856	-10.686	.000
Lag3 db	-.1631221	.0139074	-11.729	.000
Lag4 db	-.0939637	.0140831	-6.672	.000
Lag5 db	-.0619189	.0140772	-4.399	.000
Lag6 db	-.1211814	.0141262	-8.578	.000
db12	-.067737	.0084264	-8.039	.000
db24	-.0103636	.0070925	-1.461	.144
db36	-.0191648	.0065055	-2.946	.003
Lag1 bde	-.0003716	.0000187	-19.894	.000
da6	.0162051	.0072949	2.221	.026
da12	.029355	.0076464	3.839	.000
da24	.0231173	.006233	3.709	.000
da36	.0232844	.0060726	3.834	.000
bv12	.0001046	.0000154	6.778	.000
av12	6.26e-06	.0000158	.397	.692
_cons	-1.06e-06	4.26e-06	-.248	.804

B. Ask Reaction to Changes in Bid Quotes

Number of observations = 4,975
 $F(16,4959) = 51.6$
 Prob. > $F = .000$
 $R^2 = .142$
 Adjusted $R^2 = .140$
 Root MSE = .0002

Change in Ask Quote (da)	Coefficient	SE	<i>t</i>	$P > t $
Lag1 da	-.114878	.013517	-8.499	.000
Lag2 da	-.088986	.0134761	-6.603	.000
Lag3 da	-.0702402	.0135679	-5.177	.000
Lag4 da	-.0350031	.0136022	-2.573	.010
Lag5 da	-.1631024	.0136826	-11.920	.000
Lag6 da	-.0421156	.0138858	-3.033	.002
da12	-.0055925	.0074358	-.752	.452
da24	-.0156702	.0059864	-2.618	.009
da36	-.0039288	.0058462	-.672	.502
Lag1 ade	.0003776	.0000163	23.221	.000
db6	-.0027777	.0075712	-.367	.714
db12	.0160504	.0079987	2.007	.045

(continued)

Table 4.21 (continued)

Change in Ask Quote (da)	Coefficient	SE	<i>t</i>	<i>P</i> > <i>t</i>
<i>db24</i>	.0206124	.0068112	3.026	.002
<i>db36</i>	.0064569	.0062616	1.031	.303
<i>av12</i>	-.0000359	.0000152	-2.361	.018
<i>bv12</i>	.0000147	.0000148	.997	.319
<i>_cons</i>	-6.30e-06	4.08e-06	-1.542	.123

with a significant negative sign (and, even more occasionally, AD or ADQ lagged would enter with a positive sign), suggesting that stronger deal activity at the bid (ask) caused bid quotes to be lowered (raised). The same feature also occurs weakly for DA, with AD entering positively, as shown in panel *b* of table 4.20.

Again, we examined how the results would change if we ran the regressions omitting all zero price change entries (80 percent of the sample). The results are shown in table 4.22. Our process of trying to eliminate insignificant variables resulted in almost identical "best" equations, with and without zero price changes, but the relative importance of the coefficients as measured by their *t*-values changed.²⁷ The fit, as before, improves sharply once zero price changes are omitted, with the adjusted *R*² improving threefold in the bid price equation (to 0.43) and more than doubling (to 0.33) in the ask price equation. But, with a more variable series, the root MSE also again doubles.

We next compared our results for the determination of quote revision over D2000-2 with a similar exercise for FXFX (see table 4.23). The results for DFXB and DFYA showed similar features for the lagged dependent variable with strong negative autocorrelation (a first-order negative moving average pattern) and a significant role for the spread (positive in the bid equation, negative in the ask). Again as in the D2000-2 equations, volatility variables appear to enter, but in rather a complicated way. Thus, the absolute change in the *ask* price enters the determination of the change in both the ask and the bid price at two separate lags with reversed signs. Tests over a longer run of data are needed to resolve whether, and how, prior volatility affects price quote revision.

27. The absolute size of the coefficients on the lagged dependent variables increased by a factor of over three times for bid prices but nearer eight times for ask prices. With their standard errors rising by a factor of over five times in both cases, the *t*-values of the lagged dependent variables fell for bid price changes but rose for ask price changes (relative to those in table 4.20). As before, the *t*-values of the deal exhaust variables fell from nearly 20 to about 5. By contrast, the coefficient on the lagged spread variable rose sharply in the bid price equation, where the size of the coefficient rose by a factor of ten and the *t*-value also increased. (Note that we did test that the spread with six lags entered more strongly than the spread lagged once in the ask price change equation.) Otherwise, the residual variables that entered significantly changed around slightly; a variety of volatility variables still entered weakly without any clear, or intuitive, direction of effect, and, again, the effects of previous large quantities of ask deals (AD6 and ADQ24) tended to raise both bid and ask quotes.

Table 4.22 The Determinants of Quote Revision: Zero Changes Omitted: Tick by Tick

A. Bid Quote Revision				
Number of observations = 727				
$F(14,713) = 40.6$				
Prob. > $F = .000$				
$R^2 = .444$				
Adjusted $R^2 = .433$				
Root MSE = .0004				
Change in Bid Quote (db)	Coefficient	SE	t	$P > t $
Lag1 db	-.2573243	.0662203	-3.886	.000
Lag2 db	-.4178523	.0650541	-6.423	.000
Lag3 db	-.4402707	.0644304	-6.833	.000
Lag4 db	-.253194	.0779863	-3.247	.001
Lag5 db	-.0483614	.0720124	-.672	.502
Lag6 db	-.3135134	.0784248	-3.998	.000
Lag1 bde	-.0002782	.000057	-4.881	.000
db12	-.1727642	.0426325	-4.052	.000
bd6	-.0000829	.0000325	-2.554	.011
Lag1 s	.3735075	.0334393	11.170	.000
adb12	-.1342426	.0479387	-2.800	.005
adq24	.0000298	9.69e-06	3.075	.002
ada6	-.0906233	.0417378	-2.171	.030
ada24	.1076248	.0440883	2.441	.015
_cons	-.000137	.0000306	-4.471	.000
B. Ask Quote Revision				
Number of observations = 747				
$F(11,736) = 33.7$				
Prob. > $F = .000$				
$R^2 = .335$				
Adjusted $R^2 = .325$				
Root MSE = .0004				
Change in Ask Quote (da)	Coefficient	SE	t	$P > t $
Lag1 da	-.7993277	.0805063	-9.929	.000
Lag2 da	-.4331602	.0587267	-7.376	.000
Lag3 da	-.4928082	.076277	-6.461	.000
Lag4 da	-.422294	.0766422	-5.510	.000
Lag5 da	-.5945421	.0571305	-10.407	.000
Lag6 da	-.1588786	.0634126	-2.505	.012
Lag1 ade	.0003121	.0000558	5.595	.000
adq24	.0000411	8.78e-06	4.684	.000
ad6	.0000926	.0000305	3.040	.002
Lag6 s	-.125264	.0230924	-5.424	.000
adb6	.1099008	.0436925	2.515	.012
_cons	-.0000683	.0000277	-2.466	.014

Table 4.23 The Determination of Quote Changes over FXXF

A. Bid Prices				
Number of observations = 4,983				
$F(10,4973) = 96.1$				
Prob. > $F = .000$				
$R^2 = .162$				
Adjusted $R^2 = .160$				
Root MSE = .0002				
Change in FX Bid Quote (dfxb)	Coefficient	SE	t	$P > t $
Lag1 dfxb	-.3652674	.0140455	-26.006	.000
Lag2 dfxb	-.298141	.0148757	-20.042	.000
Lag3 dfxb	-.2499072	.0153796	-16.249	.000
Lag4 dfxb	-.1231102	.0155578	-7.913	.000
Lag5 dfxb	-.0775751	.0154255	-5.029	.000
Lag6 dfxb	-.05442	.015174	-3.586	.000
fxb12	-.0198112	.0096738	-2.048	.041
Lag6 s	.2664089	.0266012	10.015	.000
adfxa6	-.0256695	.011956	-2.147	.032
adfxa24	.0275647	.0097646	2.823	.005
_cons	.0001859	.000019	-9.770	.000
B. Ask Prices				
Number of observations = 4,983				
$F(9,4974) = 117.5$				
Prob. > $F = .000$				
$R^2 = .175$				
Adjusted $R^2 = .173$				
Root MSE = .0002				
Change in FX Ask Quote (dfxa)	Coefficient	SE	t	$P > t $
Lag1 dfxa	-.3707797	.0141493	-26.205	.000
Lag2 dfxa	-.3252713	.0149596	-21.743	.000
Lag3 dfxa	-.2448494	.0155451	-15.751	.000
Lag4 dfxa	-.1190115	.0155936	-7.632	.000
Lag5 dfxa	-.0820493	.015018	-5.463	.000
Lag6 dfxa	-.0497289	.0142814	-3.482	.000
adfxa6	-.0317825	.0130013	-2.445	.015
adfxa24	.0523204	.0105903	4.940	.000
Lag6 s	-.2650965	.0292663	-9.058	.000
_cons	.0001846	.0000209	8.831	.000

sion, either over D2000-2 or over FXXF. The other variables tested (i.e., the prior frequency of quote revision, the absolute change in lagged bid prices, etc.) were not significant.

In D2000-2, unlike FXXF, changes in the bid (ask) price initially become incorporated into the spread, which is highly positively correlated. Indeed, the

first-order autocorrelation with the spread in the previous five-second period has a coefficient of about 0.88 and a t -value in excess of 50, as will be shown below. In order to lessen the power of this relation and show the effects of other variables, we mostly worked with a lagged dependent variable with a thirty-second lag, Lag 6 s. Once again, a deal that exhausts a quote will force a price revision and an increase in the spread as the price shifts to the next limit order, so BDE_{t-1} and ADE_{t-1} were always entered. Thus, the basic equation was

$$S = .000220 + 0.620 S_{t-6} + .000179 BDE_{t-1} + .000313 ADE_{t-1}$$

$$(0.000015) \quad (0.011) \quad (0.000047) \quad (0.000042)$$

$$R^2 = 0.398.$$

As earlier noted, an increase in the bid price will reduce the spread, and an increase in the ask price will increase it. These results came through strongly in the equations. The standard finding is that volatility will increase spreads, and this was also strongly supported, as shown by the significant t -values on AV and BV. Our basic equation, using S_{t-6} as the lagged dependent variable, is shown in table 4.23. When S_{t-1} is introduced instead, the fit improves, but the significance of all the other variables weakens dramatically, and even the sign of the other independent variables often goes wrong since almost all their influence is incorporated into S_{t-1} , as shown in panel B of table 4.24. Besides the exhaust dummies, price revision, and volatility variables, we also looked to see whether either the event or size of deals or the frequency of quote revisions affected the spread. The answer is generally no, once the significant variables above are also entered. As can be seen from table 4.23, the number of bid deals in the thirty seconds from $t-30$ to $t-60$ (i.e., BD12) enters with a negative significant coefficient.

There is some uncertainty in the literature about what relation to expect between the volume (number) of transactions and the spread. On theoretical grounds, Admati and Pfleiderer (1988) and Foster and Viswanathan (1990) expect liquidity trading to cluster together so that low adverse selection trading costs should occur at times of high volume; yet there is evidence in both the NYSE (Foster and Viswanathan 1993) and the foreign exchange market (Glassman 1987) that the intraday pattern is for spreads to be positively correlated with volume. Bessembinder (1994) seeks to resolve this conflict by distinguishing between expected and unexpected volumes, with the expected signs on these being found to be, as hypothesized, negative and positive. We do not, however, feel that our relatively weak finding of a negative coefficient on a volume variable helps resolve this problem; we are inclined to dismiss this finding as possibly occurring by chance; its significance, along with that of many other variables, was cut back sharply when S_{t-1} was entered as the lagged dependant variable.

By contrast, there is no uncertainty in the literature that information asym-

Table 4.24 Spreads

A. With Lagged Dependent S_{t-6}

Number of observations = 4,980
 $F(16,4964) = 953.3$
 Prob. > $F = .000$
 $R^2 = .754$
 Adjusted $R^2 = .753$
 Root MSE = .0003

s	Coefficient	SE	t	$P > t $
Lag6 s	.4768366	.0120753	39.489	.000
Lag1 bde	.0002891	.0000298	9.691	.000
Lag1 ade	.0003563	.000027	13.184	.000
bv12	.0002113	.0000277	7.634	.000
bv60	.0002813	.0000269	10.450	.000
av12	.0001099	.000028	3.919	.000
av60	.0003531	.000024	14.688	.000
bd12	-.0000393	.0000107	-3.679	.000
db6	-.672806	.0127785	-52.652	.000
db12	-.2537867	.0154748	-16.400	.000
db24	-.1797322	.0128506	-13.986	.000
db36	-.0684526	.0107952	-6.341	.000
da6	.7448953	.0117689	63.293	.000
da12	.3830915	.0148889	25.730	.000
da24	.2522268	.0118793	21.232	.000
da36	.1455865	.0104181	13.974	.000
_cons	.0001214	8.33e-06	14.578	.000

B. With Lagged Dependent, S_{t-1}

Number of observations = 4,980
 $F(16,4964) = 1460.7$
 Prob. > $F = .000$
 $R^2 = .824$
 Adjusted $R^2 = .824$
 Root MSE = .0003

s	Coefficient	SE	t	$P > t $
Lag1 s	.8807202	.0136244	64.643	.000
Lag1 bde	.0003409	.0000252	13.517	.000
Lag1 ade	.0003596	.0000228	15.750	.000
bv12	-.0000139	.0000239	-.582	.560
bv60	.000057	.0000233	2.449	.014
av12	.0000297	.0000237	1.252	.210
av60	.0000282	.0000216	1.310	.190
bd12	-8.43e-06	9.02e-06	-.934	.350
db6	.042907	.0150085	2.859	.004
db12	.0280221	.0144781	1.935	.053
db24	-.0019636	.0115459	-.170	.865
db36	.0107807	.0092818	1.161	.245
da6	-.0089642	.0153301	-.585	.559

Table 4.24 (continued)

s	Coefficient	SE	t	P > t
da12	.0435953	.0146797	2.970	.003
da24	.021478	.0112872	1.903	.057
da36	.009385	.0093128	1.008	.314
_cons	.00003	7.33e-06	4.099	.000

metries and high volatility will be associated with high spreads.²⁸ This has been found in two recent articles using FXXF data. Bollerslev and Melvin (1994) and Bessembinder (1994). We have, however, shown earlier (tables 4.10 and 4.11 above) that the form of the (numerical) relation (the coefficients) between volatility and spreads differs depending on whether D2000-2 or FXXF data are used.

So next, for comparison, we examined the determination of spreads on FXXF for the same deutsche mark/dollar exchange rate over the same period. The results of this (see table 4.25) show that, besides positive autocorrelation (although much weaker than in D2000-2, the coefficient on the first lag drops from 0.88 to 0.38), the spread is again positively related to volatility (ADFXB24). There is also a weak relation with the frequency of quote entry, but the coefficients are of equal and opposite sign, so the net effect is negligible. Most of the variation in spreads in FXXF is just noise, with an adjusted R^2 of 0.15, as compared with over 0.75 for D2000-2.

We then looked at the factors affecting the absolute change in prices (a measure of the volatility) of bid (ADB) and ask quotes (ADA) both in D2000-2 and in FXXF. The results of this part of the exercise were not particularly exciting (see tables 4.26 and 4.27 as well as n. 29).

28. Much of the literature on spreads, especially for spreads in the NYSE, seeks to distinguish between the effects of trading costs, inventory costs, and information asymmetry (e.g., Madhavan and Smidt 1991). We cannot attempt a similar exercise as we have no measure of inventories, unlike Lyons (1995).

29. Obviously, the exhaustion of the quote by a deal would cause a jump in prices, so, in the equations to explain the absolute change in prices in D2000-2, ADB and ADA, BDE and ADE were entered into their respective equations. The lagged dependent variable and the absolute change in quote revision on the other side (e.g., ADA in the ADB equation) were quite strongly significant. The prior event of deals (AD and BD), but *not* their size (BDQ and ADQ), and the frequency of price revision (AF and BF) were also a positive, but somewhat weaker, influence on the absolute value (volatility) of price change. The size of ask quotes (AQ) appeared to affect the absolute value of ask price changes, although the two lags that entered had offsetting effects. Two of our (better) representative equations are given in table 4.26. Again, we undertook the companion exercise of looking at absolute price changes on FXXF. Apart from the lagged dependent variables, the spread entered with a significant positive coefficient. Presumably, this is picking up some (expected) determinants of volatility (not otherwise caught by the lagged dependant variables). The change in ask prices enters the equation explaining the absolute change in ask prices, whereas the frequency of quote entries enters the similar equation for the bid prices. With price movements in the bid and ask being much more closely tied together and similar for FXXF than for D2000-2, here we show only the former equation in table 4.27 since the latter (apart from the substitution of FXBF for DFXA) is almost identical.

Table 4.25 Determination of FAFX Spreads

s	Coefficient	SE	t	P > t
Number of observations = 4,927				
$F(5,4967) = 180.3$				
Prob. > F = .000				
$R^2 = .153$				
Adjusted $R^2 = .152$				
Root MSE = .0002				
Lag1 s	.377904	.0131287	28.785	.000
Lag4 s	.0300787	.013117	2.293	.022
fxaf12	4.58e-06	2.76e-06	1.660	.097
fxaf36	-4.94e-06	1.76e-06	-2.803	.005
adfb24	.0353665	.0089404	3.956	.000
_cons	.0004243	.0000179	23.654	.000

As described earlier, the frequency of quote revision (BF and AF) Granger-causes the event of deals. The reverse causal relation also holds, with the number of recent deals influencing the frequency of quote revision. This is consistent with the hypothesis that trading activity itself generates revisions of prior information and hence further trading (e.g., French and Roll 1986). Thus, BD6 is the dominant influence on BF and AD6 on AF. Besides this, there is a weak positive effect from the lagged dependent variable and from the lagged frequency of the other quote (AF in the equation for BF, and vice versa), some positive effect of higher price volatility on the frequency of quote revision, and, finally, a weak and rather uncertain (the lagged variables usually had an offsetting effect) effect from the quote size variables (BQ and AQ). We show two of our better representative equations in table 4.28.³⁰

Once again, largely for the record, we ran associated regressions to examine the determinants of the frequency of quote entry over FAFX. This showed that, apart from own lagged values, the only variable, from the set of FAFX data available examined here, that influenced the frequency of quote entry over FAFX was a lagged volatility variable.³¹ In order to save space, the table is not shown but is available from the authors on request.

Finally, in this set of studies of activity on D2000-2 (and FAFX), we explored the determinants of the quantities posted, BQ and AQ. (Recall that we chose *not* to seek to examine the determinants of the size of deal, BDQ and ADQ, since these most often just exhausted the quantity already on offer.) A noteworthy feature of our results is that the quantities posted, BQ and AQ, did not significantly affect most of the preceding variables (e.g., probability of

30. We have no good explanation for the negative values for AD24 or ADA24 in the equation shown in panel A of table 4.28, and we would again be inclined to regard these as chance findings.

31. This volatility variable was the absolute change in prices over the preceding half minute (ADFXA in the ask equation and ADFXB in the bid equation).

Table 4.26 The Determinants of Absolute Price Changes

A. In Bid Prices (adb)				
Number of observations = 4,980				
$F(7,4973) = 78.3$				
Prob. $> F = .000$				
$R^2 = .099$				
Adjusted $R^2 = .098$				
Root MSE = .0002				
In Bid Prices (adb)	Coefficient	SE	t	$P > t $
adb6	.0638526	.0082624	7.728	.000
adb12	.0190271	.0083937	2.267	.023
abd24	.0366817	.0073459	4.994	.000
Lag1 bde	.0003333	.0000181	18.460	.000
ada12	.0178956	.0078319	2.285	.022
ada36	.0156567	.0064971	2.410	.016
af24	6.13e-06	2.02e-06	3.037	.002
_cons	3.55e-06	5.04e-06	.705	.481
B. In Ask Prices (ada)				
Number of observations = 4,980				
$F(11,4969) = 82.2$				
Prob. $> F = .000$				
$R^2 = .154$				
Adjusted $R^2 = .152$				
Root MSE = .0002				
In Ask Prices (ada)	Coefficient	SE	t	$P > t $
ada6	.0621732	.0079386	7.832	.000
ada12	.0244177	.0076472	3.193	.001
ada24	.0184217	.0063198	2.915	.004
ad24	.0000145	3.37e-06	4.308	.000
ad36	.0000113	3.35e-06	3.392	.000
ade1	.0003508	.0000155	22.667	.000
af6	5.86e-06	3.17e-06	1.851	.064
adb36	.0210127	.0069293	3.032	.002
bf36	5.46e-06	2.00e-06	2.726	.006
aq12	8.91e-07	3.79e-07	2.348	.019
aq24	-7.22e-07	2.14e-07	-3.365	.000
_cons	-6.84e-06	7.87e-06	-.869	.385

deal, quote revision, spread) and only weakly affected, if at all, volatility and the frequency of quote entry. Anyhow, the main factors affecting the quantities offered, BQ and AQ, as shown in table 4.29, are the respective lagged dependent variables, with strongly significant first-order positive autocorrelation (but in the case of BQ thereafter a somewhat complex dynamic process), and the number of prior deals (BD in the BQ equation, AD in the AQ equation), which

Table 4.27 The Determinants of Absolute Price Changes on FFXF (adfxa)

Determinants of Absolute Price Changes on FFXF (adfxa)				
	Coefficient	SE	<i>t</i>	<i>P</i> > <i>t</i>
Number of observations = 4,972				
$F(5,4966) = 51.7$				
Prob. > $F = .000$				
$R^2 = .049$				
Adjusted $R^2 = .048$				
Root MSE = .0002				
Determinants of Absolute Price Changes on FFXF (adfxa)	Coefficient	SE	<i>t</i>	<i>P</i> > <i>t</i>
adfxa6	.1366876	.0117699	11.613	.000
adfxa12	.0338006	.011855	2.851	.004
adfxa36	.0267606	.0095021	2.816	.005
s6	.0990903	.0258231	3.837	.000
dfxa6	.0375036	.008251	4.545	.000
_cons	.0000414	.0000187	2.220	.026

reduces quote size. Other activity variables, such as BF, AF, and BDQ, enter weakly and often with offsetting signs, so their net effect is negligible. A volatility variable (BV12) enters the BQ equation positively. The only factors, however, about which we have some confidence are those for the lagged dependent variable and the negative effect of deal activity on quote size.

This extended series of results and tables must seem quite complicated, and so in a manner it is. We try to simplify by illustrating, in figures 4.14a–c, the main interrelations (*excluding* interactions whereby bid variables affect ask variables, and vice versa), with the direction of causality given by the arrow, the strong relationships displayed in figure 4.14a, the weak relationships in figure 4.14b, and the questionable relationships in figure 4.14c. A key point is that deals mainly affect quote (price) revisions, spreads, and volatility if they have exhausted the amount then on offer, but with a much weaker effect otherwise. This deal exhaustion effect is the main link *from* the deal occurrence/frequency of quote revision nexus (one-way) *to* volatility and to the quote revision/spread nexus.

The exercises, whose results were reported in these figures, were mostly, except for tables 4.19 and 4.22, done on a clock-time scale. We were both encouraged and slightly surprised to find that, when we changed the periodicity (table 4.17 compared with table 4.15) or the scale (table 4.18 compared with table 4.19 and table 4.20 compared with table 4.22), the patterns of the basic relations, as measured by the *t*-values on the key variables, remained quite robust.

4.4.3 Conditional Heteroskedasticity in D2000-2

Most asset price series exhibit ARCH, autoregressive conditional heteroskedasticity. We next turned to examine whether our price series, DB and DA,

Table 4.28 The Frequency of Quote Entry on D2000-2

A. Of Bid Prices (bf)				
Number of observations = 4,980				
$F(11,4969) = 26.6$				
Prob. > $F = .000$				
$R^2 = .055$				
Adjusted $R^2 = .053$				
Root MSE = .3435				
Of Bid Prices (bf)	Coefficient	SE	t	$P > t $
bf6	-.0093551	.0059631	-1.569	.117
bf12	-.0000455	.0053486	-.008	.993
bf24	.0109553	.0034244	3.199	.001
bf36	.0078274	.0033847	2.313	.021
bd6	.1440359	.0111967	12.864	.000
af24	.0149482	.0037679	3.967	.000
ad24	-.0157613	.0062909	-2.505	.012
ada24	-24.13405	10.76621	-2.242	.025
bv12	.0900122	.0254729	3.534	.000
bq24	.001001	.0003442	2.908	.004
bq36	-.0008411	.0003457	-2.433	.015
_cons	.0607582	.0143452	4.235	.000
B. Of Ask Prices (af)				
Number of observations = 4,980				
$F(8,4972) = 32.2$				
Prob. > $F = .000$				
$R^2 = .049$				
Adjusted $R^2 = .047$				
Root MSE = .3484				
Of Ask Prices (af)	Coefficient	SE	t	$P > t $
af6	.0166517	.005561	2.994	.003
ad6	.0963283	.0092722	10.389	.000
bf12	.0118653	.0051196	2.318	.021
bf36	.0109036	.0033706	3.235	.001
adq24	.0049921	.0026929	1.854	.064
adq36	.0070069	.0027153	2.581	.010
av60	.0510773	.0199153	2.565	.010
aq24	-.0006868	.0003179	-2.160	.031
_cons	.0705923	.0133874	5.273	.000

also had such characteristics, either in clock (five-second) time or on a tick-by-tick (activity) scale. We could also explore whether the addition of transaction data (e.g., BD, BDE) would influence the GARCH coefficients. Having already examined the relation between the GARCH coefficients of the interpolated D2000-2 and FFX series in section 4.3, we now focus solely on the former to investigate whether, in clock time or using a data set constructed solely using quote and transaction activity, the series exhibit signs of condi-

Table 4.29 The Determinants of Quote Size

A. Bid Quote Size (bq)				
Number of observations = 4,980				
$F(9,4971) = 329.2$				
Prob. > $F = .000$				
$R^2 = .373$				
Adjusted $R^2 = .372$				
Root MSE = 1.393				
Bid Quote Size (bq)	Coefficient	SE	t	$P > t $
bq6	.1505502	.0032501	46.321	.000
bq12	-.0257189	.0030404	-8.459	.000
bq24	.0081696	.0014368	5.686	.000
bd6	-.1266831	.0414378	-3.057	.002
bf12	-.0550551	.0215234	-2.558	.011
bf24	.0343499	.0134897	2.546	.011
bv12	.217296	.0993405	2.187	.029
af24	-.0533222	.0133771	-3.986	.000
af36	.0394126	.0135053	2.918	.004
_cons	.6453671	.058449	11.042	.000
B. Ask Quote Size (aq)				
Number of observations = 4,981				
$F(11,4970) = 235.6$				
Prob. > $F = .000$				
$R^2 = .342$				
Adjusted $R^2 = .341$				
Root MSE = 1.458				
Ask Quote Size (aq)	Coefficient	SE	t	$P > t $
aq6	.1242625	.0029287	42.430	.000
aq24	.0095996	.0014215	6.753	.000
ad6	-.1941901	.0391663	-4.958	.000
ad24	-.0502285	.0234261	-2.144	.032
af6	.0641108	.0231159	2.773	.006
bf6	-.0892569	.0230461	-3.873	.000
bdq6	.053556	.0222602	2.406	.016
bdq12	.0597338	.0208646	2.863	.004
bdq36	-.0569668	.0140248	-4.062	.000
bq6	.0089876	.0034053	2.639	.008
bq12	-.0097294	.0029193	-3.333	.000
_cons	.7090902	.0635584	11.157	.000

tional heteroskedasticity. The basic specification that we used is shown below. Quote revisions are assumed to depend on their own first lag and a dummy indicating a deal that exhausted the quantity on offer at the prevailing price in the previous period. The volatility expression is based on a simple GARCH(1,1), extended subsequently to examine the effect of deals on volatility:




Fig. 4.14a Strong relationships: main transmission channels




Fig. 4.14b Weak relationships: but some clear effect




Fig. 4.14c Questionable relationships

$$\Delta b_t = \alpha_0 + \alpha_1 \Delta b_{t-1} + \alpha_2 \text{BDE}_{t-1} + \varepsilon_t, \quad \varepsilon_t | I_{t-1} \sim N(0, h_t),$$

$$h_t = \beta_0 + \beta_1 \varepsilon_{t-1}^2 + \beta_2 h_{t-1} + \beta_3 \text{BDE}_{t-1} + \beta_4 \text{BD}_{t-1}.$$

For brevity's sake, we report the results only for the bid side of the deutsche mark/dollar, presented in tables 4.30 and 4.31.

Taking first the estimations for the quote revision equations, note that the autoregressive parameter, α_1 , is negative in all cases but with a consistently greater magnitude for the activity scale data, as previously reported when comparing tables 4.18 and 4.19 and tables 4.20 and 4.22. As before, the deal exhaust indicator has the expected negative effect on quote revisions.

Then, next inspecting the volatility estimations, GARCH effects are present in both data sets. As shown in tables 4.30 and 4.31, the parameters β_1 and β_2 are significantly different from zero in a standard GARCH(1,1). We then examined whether deals affected quote revisions in an indirect manner through the underlying volatility series. This was done by adding lagged deal and deal exhaust indicators to the simple GARCH framework. For our activity scale data, we could not uncover any real effect of deals on volatility. Neither of the previously defined variables, BD_{t-1} and BDE_{t-1} , entered significantly into our

Table 4.30 GARCH Estimation, Including Transactions, Calendar Time Data

	α_0	α_1	α_2	β_0	β_1	β_2	β_3	β_4
GARCH	1.2e-5 ^a	-.166 ^a	-2.2e-4 ^a	1.2e-9 ^c	.291 ^a	.874 ^a
+bde _{t-1}	1.4e-5 ^a	-.144 ^a	-2.6e-4 ^a	5.7e-9 ^b	.113 ^a	.734 ^a	2.4e-8 ^b	...
+bd _{t-1}	1.3e-5 ^a	-.148 ^a	-2.8e-4 ^a	7.1e-9 ^a	.213 ^a	.667 ^a	...	2.7e8 ^b

Note: Presentation of the estimated parameters of the specification described in section 4.4.3 for the five-second data set plus extended specifications including lagged deal and deal exhaust indicators in the volatility expression.

^aSignificantly different from zero at 1 percent.

^bSignificantly different from zero at 5 percent.

^cInsignificantly different from zero.

Table 4.31 GARCH Estimations Including Transactions, Tick-by-Tick Data

	α_0	α_1	α_2	β_0	β_1	β_2	β_3	β_4
GARCH	5.4e-5 ^a	-.436 ^a	-3.3e-4 ^a	1.1e-8 ^a	.332 ^a	.681 ^a
+bde _{t-1}	5.4e-5 ^a	-.423 ^a	-3.1e-4 ^a	1.3e-8 ^c	.252 ^a	.724 ^a	-1.2e-8 ^c	...
+bd _{t-1}	5.5e-5 ^a	-.433 ^a	-3.1e-4 ^a	1.4e-8 ^c	.241 ^a	.730 ^a	...	-1.2e-8 ^c

Note: Presentation of the estimated parameters of the specification described in section 4.4.3 for the activity-based data set plus extended specifications including lagged deal and deal exhaust indicators in the volatility expression.

^aSignificantly different from zero at 1 percent.

^bSignificantly different from zero at 5 percent.

^cInsignificantly different from zero.

estimation, and, indeed, their negative sign is implausible. But, when we moved to the five-second data set, the results were markedly different. Both of these variables had a positive effect on volatility, significant at the 5 percent level. Maybe when the quiet no-change observations are excluded in the activity-based data, the slighter persistent effects of deals on volatility become drowned out in the noisier "news."

Indeed, in all cases, the GARCH estimates are far more significant in the five-second, clock-time data set. Comparison of the *t*-statistics between the activity and the five-second data shows those in the latter case to be far greater. This does not, however, imply that GARCH-type phenomena are *better* addressed in clock time than in activity time. It has been suggested that GARCH effects apparent in clock-time data may be the result of the transformation of a uniform, latent process that evolves on a different (activity) scale (see Stock 1988). This could underlie the diminished significance of the GARCH parameters in the tick-time results. This, however, is a subject for further research and is not pursued further here.

4.4.4 A Comparison with Hasbrouck's (1991) NYSE Study

Finally, in his 1991 study of the NYSE, Hasbrouck studied a bivariate VAR of the interrelations between deals (and/or deal quantities) and price revision (taking the middle of the bid-ask quote), scaling by activity, tick time. Here, we show his main results (which he gives in his table II on p. 194) and our replication from our own data, both in tick time as he ran the regressions and in clock time, here reported in table 4.32.

Since the scales of the price changes in the two markets (NYSE and foreign exchange) are markedly different, the differences between the absolute sizes of the coefficients should be ignored and are not shown (but are available from the authors). What matters is the size and pattern of the *t*-values, as shown in table 4.32. This shows that the equation for price quote revisions, the *a* and *b* *t*-values in the first columns, are qualitatively similar. In both cases, although considerably more strongly in the foreign exchange data, both in the clock-time and in the activity-time equations, there is significant negative autocorrelation, and in both cases quote revisions are strongly positively related to prior deals (i.e., a sell causes a drop in prices and a buy an appreciation). Like several other economists, Hasbrouck tended to dismiss the negative autocorrelation, noting that it "may simply arise from measurement error" (1991, 195); our repeated findings of such negative autocorrelation on high-frequency foreign exchange data make us believe that this finding cannot be brushed aside in this fashion. In the activity-based foreign exchange equation (column B), we can explain considerably more of the fluctuations in quote revisions than Hasbrouck, but this is primarily due to the stronger autocorrelation. (Hasbrouck does not report the *F*-statistic showing the combined effect of the x_0 variable on r [in our case it is 10.25], but a look at the comparable *t*-statistics suggests that the combined effect may be somewhat stronger in his equations.)

The main, qualitative difference between his results and our own comes in

Table 4.32 Estimates of the Bivariate Vector Autoregressive Model

	A	B	C		A	B	C
a_1	-7.22	-16.16	-8.19	c_1	-13.44	-1.47	-1.33
a_2	-.67	-7.38	-7.09	c_2	-6.05	-.07	.94
a_3	-.17	-5.36	-7.22	c_3	-1.80	-.72	.59
a_4	-1.31	-6.39	-1.65	c_4	-.46	.57	-.25
a_5	-.14	-3.00	-5.65	c_5	.41	1.21	.79
b_0	15.15	7.57	-.69	d_1	10.16	4.82	2.07
b_1	6.83	2.87	13.53	d_2	7.20	2.45	2.53
b_2	.46	3.09	.42	d_3	4.66	4.15	1.25
b_3	.87	.13	1.87	d_4	1.24	.67	3.73
b_4	-.30	2.61	.69	d_5	2.03	.87	1.85
b_5	.94	2.36	3.55	R^2	.085	.038	.005
R^2	.096	.175	.068				

Note: We estimate a five-lag near-VAR involving r_t , the revision in the quote midpoint, and x_{0t} , the trade indicator variable. The VAR is not exact as the trade indicator is assumed to have a contemporaneous effect on quote revisions, as shown in the system below:

$$r_t = \sum_{i=1}^5 a_i r_{t-i} + \sum_{i=0}^5 b_i x_{0t-i} + v_{1t}, \quad x_{0t} = \sum_{i=1}^5 c_i r_{t-i} + \sum_{i=1}^5 d_i x_{0t-i} + v_{2t}.$$

t-statistics are reported for each of the estimated parameters. Column A reproduces Hasbrouck's results, column B gives our equivalent activity-scale results, and column C gives our results on a clock-time basis.

the second set of equations for the event of deals. In Hasbrouck's equation, price quote revisions have a significant negative effect on deals; the first two c coefficients have *t*-values well below -2 . In our own work, as reported in tables 4.15 and 4.16 above, price quote revisions have little, or no, effect on the probability of deals occurring, and this (negative) finding recurs also here. Both in Hasbrouck's results and in our own, there is positive autocorrelation in deal events, slightly stronger in his case than in either of our two runs. So, although the fit in all cases is close to zero, Hasbrouck can "explain" rather more of deal eventuality than we can. Hasbrouck notes that "a negative relation between trades and lagged quote revisions is consistent with inventory control effects since a monopolistic marketmaker with an inventory surplus would reduce his quotes to elicit more purchases. It is also consistent with the price experimentation hypothesis of Leach and Madhavan (1989) in which the marketmaker sets quotes to extract information optimally from the traders. These possibilities are deserving of further study" (1991, 295). In this further study, we find that, in our data sample from a market with many competing marketmakers, there was no indication of any significant (negative) effect between trades and lagged quote revisions.

In addition, we examined whether our results were robust to a longer lag structure (ten instead of five); the answer was yes. We were also able to replicate with our data the exercises done by Hasbrouck (1991) in his tables III (p. 198) and V (p. 203). In table III, Hasbrouck examines the interrelations be-

Table 4.33 Comparison of Coefficients

	Own Lags	x_0	x	x^2
Hasbrouck	32.29	.98	5.89	-3.24
Our data	138.69	2.72	-.018	.118

tween r , price quote revisions, x_0 , the event of a deal, x , the size of a deal (— for sales, + for purchases), and $x^2 = x_0 \cdot |x|^2$. Our results (the data are available on request from the authors) show generally less effect of the deal (x) variables on quote revisions. Unlike in Hasbrouck, neither the x nor the x^2 variables have a significant effect on r ; only the x_0 variables do. Like Hasbrouck, we find that none of the variables, even the lagged dependent variables, can help much to explain x , the size of deals; indeed, as in his study, the lagged event of a deal x_0 is very slightly better at explaining x , the size of deals, than lags of x themselves. Again, as in Hasbrouck's work, neither the size nor the squared size of deals, x and x^2 , has any effect on the eventuality of deals, x_0 —indeed, even less in our data than in his.

Finally, we look at the determinants of the spread. In his table V (1991, 203), Hasbrouck regresses the spread, for his particular equity share, Ames Department Stores, on its own (five) lags and the absolute values of the current and five lagged values of x_0 , x , and x^2 . The t -statistics for the *sums* of these variables—and for comparison on our data (activity scale only)—are as shown in table 4.33. This shows that the extent of the positive autocorrelation of spreads is even larger in our data than in his. Otherwise, the significance of deals in our data set is rather less than in his and works in our own case primarily through x_0 , the event of a deal, rather than its size (or squared size). In particular, Hasbrouck finds some general tendency for the effect of x (on quote revisions and spreads) to be positive and for x^2 to be negative, which we do not find in our data set; but this is very likely because of the manner in which deal size was limited by the usually small size of the quote on offer in our data set, as earlier described.

4.4.5 Conclusions

It is now time to summarize this long, and often quite complex, study of the interrelations and determinants of the variables that can be extracted from D2000-2: for example, event, price, and size of deal, and whether an order exhausts the prior quote; the frequency of entry, price, size, and volatility of prices for both the bid and the ask; and the spread between them. Let us do so by reviewing our main findings.

1. Unlike the price quote series, which exhibits highly significant negative autocorrelation at high frequencies, the transaction price series exhibits no strong signs of autocorrelation (there is an insignificant negative first-order

autocorrelation balanced by just significant higher-order positive autocorrelation). If one could *not* observe the “bounce” between deals at the bid and ask, the transactions series would then *appear* to exhibit weak negative first-order autocorrelation.

2. Tests of length of runs of deals at the bid and ask suggest that these have a fat (long) tail, which in this data set appears to be associated with strong price trends.

3. Studies of interactions between the many variables available from D2000-2 suggest a close interrelation (nexus) between quote frequency and deals (two-way causality) and between quote (price) changes and the spread (two-way causality). These two nexuses are linked, in that a deal that exhausts the amount offered at a previously quoted price will cause a price change both directly and indirectly via its effect on the spread (both directly and again indirectly by raising volatility). Deals that do not exhaust the amount on offer have a much weaker effect. There are only weak relations (in either direction) between the quantities (posted) and any of the other variables in the system.

4. Unlike a single dealer system, where the dealer will normally adjust both bid and ask quotes simultaneously, in this multiple competitive dealer system the bids and asks are normally input by different banks. There is no automatic reason why bid quotes should be revised in response to changes in ask quotes (or deals). In practice here, they did *not* respond much to such activity on the other side. Instead, price changes on one side primarily affected the spread and thence *gradually* the quote on the other side, with the spread acting as an error correction mechanism between the cointegrated bid and ask series.

5. The main pattern of relations reported in point 3 above appear to be encouragingly robust, as evidenced by the similarity of *t*-values, to changes in either the periodicity or the scale over which the regressions were run.

6. On the other hand, the GARCH equations varied considerably when run in clock-time rather than on an activity scale. The results for the former were more intuitive.

7. We were able to run an exact comparison, and replication, of Hasbrouck’s (1991) study of transaction/quote relations in the NYSE. The main difference between us is that in his study lagged quote revisions have a significant (negative) effect on deals, whereas there is no such interaction in our data set.

4.5 Tailpiece

We have already summarized our main findings at the ends of sections 4.3 and 4.4. Here, we wish to emphasize again how short our data period was, only seven hours. Our findings should, therefore, be treated with due caution. By the same token, there would be considerable value, not only to academics but also to practitioners, in obtaining additional data of this high-quality format. We hope, and expect, that such data will become more widely available soon.

Appendix

Reuters D2000-2 Data




Fig. 4A.1 Deutsche mark/dollar hour 0: 0-20




Fig. 4A.2 Deutsche mark/dollar hour 0: 20-40




Fig. 4A.3 Deutsche mark/dollar hour 0: 40-59




Fig. 4A.4 Deutsche mark/dollar hour 1: 0-20




Fig. 4A.5 Deutsche mark/dollar hour 1: 20-40




Fig. 4A.6 Deutsche mark/dollar hour 1: 40-59




Fig. 4A.7 Deutsche mark/dollar hour 2: 0-20




Fig. 4A.8 Deutsche mark/dollar hour 2: 20-40




Fig. 4A.9 Deutsche mark/dollar hour 2: 40-59




Fig. 4A.10 Deutsche mark/dollar hour 3: 0-20




Fig. 4A.11 Deutsche mark/dollar hour 3: 20-40




Fig. 4A.12 Deutsche mark/dollar hour 3: 40-59




Fig. 4A.13 Deutsche mark/dollar hour 4: 0-20

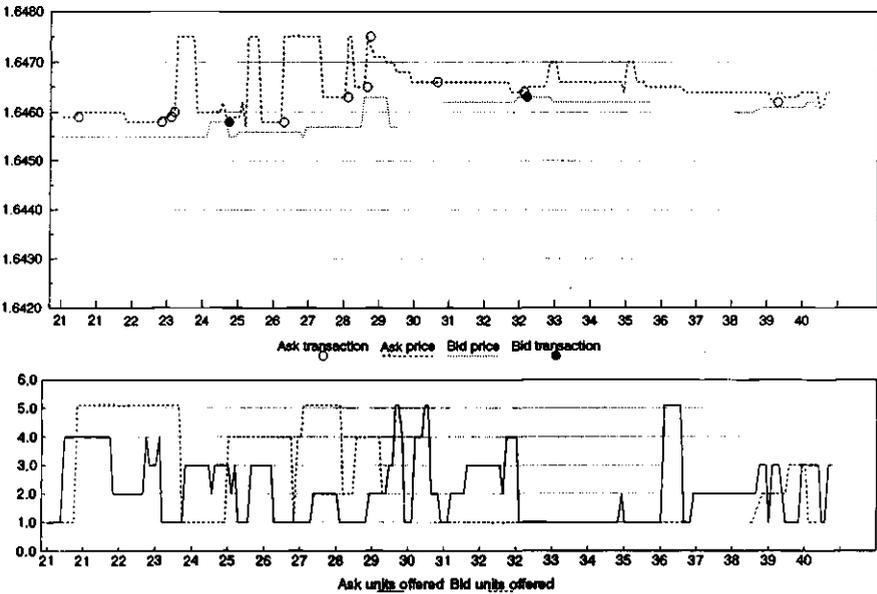


Fig. 4A.14 Deutsche mark/dollar hour 4: 20-40




Fig. 4A.15 Deutsche mark/dollar hour 4: 40-59




Fig. 4A.16 Deutsche mark/dollar hour 5: 0-20




Fig. 4A.17 Deutsche mark/dollar hour 5: 20-40




Fig. 4A.18 Deutsche mark/dollar hour 5: 40-59




Fig. 4A.19 Deutsche mark/dollar hour 6: 0-20




Fig. 4A.20 Deutsche mark/dollar hour 6: 20-40




Fig. 4A.21 Deutsche mark/dollar hour 6: 40-59




Fig. 4A.22 Deutsche mark/dollar hour 7: 0-20

References

- Admati, A., and P. Pfleiderer. 1988. A theory of intraday patterns: Volume and price variability. *Review of Financial Studies* 3:593–624.
- . 1989. Divide and conquer: A theory of intraday and day-of-the-week mean effects. *Review of Financial Studies* 2, no. 2: 189–223.
- Baillie, R. T., and T. Bollerslev. 1991. Intra-day and inter-market volatility in foreign exchange rates. *Review of Economic Studies* 58:565–85.
- Bank for International Settlements (BIS). 1993. *Survey of foreign exchange market activity*. Basle: BIS, Monetary and Economic Department.
- Bessembinder, H. 1994. Bid-ask spreads in the inter-bank foreign exchange markets. *Journal of Financial Economics* 35, no. 3 (June): 317–48.
- Blitz, J. 1993. Foreign exchange dealers enter the 21st century. *Financial Times*, 13 September 1993, 19.
- Bollerslev, T., and I. Domowitz. 1991. Price volatility, spread variability and the role of alternative market mechanisms. *Review of Futures Markets* 10:78–102.
- . 1993. Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance* 48, no. 4 (September): 1421–43.
- Bollerslev, T., and M. Melvin. 1994. Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics* 36:355–72.
- Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb. 1981. Transaction costs, order placement strategy, and existence of the bid-ask spread. *Journal of Political Economy* 89, no. 2:287–305.
- Dacorogna, M. M., U. A. Müller, R. J. Nagler, R. B. Olsen, and O. V. Pictet. 1993. A geographical model for the daily and weekly seasonal volatility in the FX market. *Journal of International Money and Finance* 12, no. 4:413–38.
- Dimson, E., ed. 1988. *Stock market anomalies*. Cambridge: Cambridge University Press.
- Domowitz, I. 1990. The mechanics of automated trade execution systems. *Journal of Financial Intermediation* 1 (June): 167–94.
- . 1993. A taxonomy of automated trade execution systems. *Journal of International Money and Finance* 12, no. 6 (December): 607–31.
- Ederington, L. H., and J. H. Lee. 1993. How markets process information: News releases and volatility. *Journal of Finance* 48, no. 4 (September): 1161–91.
- Flood, M. D. 1994. Market structure and inefficiency in the foreign exchange market. *Journal of International Money and Finance* 13, no. 2 (April): 131–58.
- Foster, F. D., and S. Viswanathan. 1990. A theory of interday variations in volumes, variances and trading costs in securities markets. *Review of Financial Studies* 3:593–624.
- . 1993. Variations in trading volume, return volatility and trading costs: Evidence on recent price formation models. *Journal of Finance* 48, no. 1 (March): 187–211.
- French, K. R., and R. Roll. 1986. Stock return variance: The arrival of information and the reaction of traders. *Journal of Financial Economics* 17:5–26.
- Glass, G. R. 1994. Multinet's FX netting solution. *Proceedings of the International Symposium on Banking and Payment Services*, 152–67. Washington, D.C.: Board of Governors of the Federal Reserve System.
- Glassman, D. 1987. Exchange rate risks and transactions costs: Evidence from bid-ask spreads. *Journal of International Money and Finance* 6:479–90.
- Goodhart, C. 1989. "News" and the foreign exchange market. Pamphlet. Manchester: Manchester Statistical Society, 17 October.

- Goodhart, C., and R. Curcio. 1991. The clustering of bid/ask prices and spreads in the foreign exchange market. Discussion Paper no. 110. Financial Markets Group, London School of Economics, January.
- Goodhart, C., and A. Demos. 1990. Reuters screen images of the foreign exchange market: The deutschemark/dollar spot rate. *Journal of International Securities Markets* 4 (Winter): 333–48.
- Goodhart, C., and A. Demos. 1991a. The Asian surprise in the forex markets. *Financial Times*, 2 September, 13.
- . 1991b. Reuters screen images of the foreign exchange market: The yen/dollar and sterling/dollar spot market. *Journal of International Securities Markets* 5 (Spring): 35–64.
- Griffiths, M. D., and R. W. White. 1993. Tax-induced trading and the turn-of-the-year anomaly: An intraday study. *Journal of Finance* 48, no. 2 (June): 575–98.
- Hasbrouck, J. 1991. Measuring the information content of stock trades. *Journal of Finance* 46, no. 1 (March): 179–207.
- Hasbrouck, J., and T. S. H. Ho. 1987. Order arrival, quote behaviour and the return-generating process. *Journal of Finance* 42, no. 4 (September): 1035–48.
- Leach, C., and A. Madhavan. 1989. Price experimentation and market structure. Working paper. Wharton School, University of Pennsylvania.
- Lease, R., R. Masulis, and J. Page. 1991. An investigation of market microstructure impacts on event study returns. *Journal of Finance* 46: 1523–36.
- Lee, C. M. C., and M. J. Ready. 1991. Inferring trade direction from intraday data. *Journal of Finance* 46:733–46.
- Lyons, R. 1993. Information intermediation in the microstructure of the foreign exchange market. Business School, University of California, Berkeley. Typescript.
- . 1995. Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics* 39:321–51.
- Madhavan, A., and S. Smidt. 1991. A Bayesian model of intraday specialist pricing. *Journal of Financial Economics* 30:99–134.
- McInish, T. H., and R. A. Wood. 1985. An analysis of transactions data for the Toronto Stock Exchange. *Journal of Banking and Finance* 14:441–58.
- . 1990. A transactions data analysis of the variability of common stock returns during 1980–1984. *Journal of Banking and Finance* 14:99–112.
- . 1991. Autocorrelation of daily index returns: Intra-day-to-day versus close-to-close intervals. *Journal of Banking and Finance* 15:193–206.
- Müller, U. A., M. M. Dacorogna, R. B. Olsen, O. Pictet, M. Schwarz, and C. Morgengegg. 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.
- Petersen, M. A., and D. Fialkowski. 1994. Posted versus effective spreads: Good prices or bad quotes. *Journal of Financial Economics* 35, no. 3 (June): 269–92.
- Pictet, O. V., M. M. Dacorogna, U. A. Müller, and C. G. De Vries. 1994. The distribution of extremal foreign exchange rate returns and extremely large data sets. Preprint. Zurich: Olsen and Associates Research Group, 22 June.
- Roll, R. 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance* 39:1127–39.
- Stock, J. 1988. Estimating continuous time processes subject to time deformation: An application to postwar U.S. GNP. *Journal of the American Statistical Association* 83, no. 401 (March): 77–85.
- Wood, R. A., T. H. McInish, and J. K. Ord. 1985. An investigation of transaction data for NYSE stocks. *Journal of Finance* 40, no. 3 (July): 723–39.
- Zhou, B. 1992. High frequency data and volatility in foreign exchange rates. Department of Finance, Sloan School of Management, Massachusetts Institute of Technology. Typescript.

Comment Richard K. Lyons

The authors do a lovely job with an important topic. The paper provides much information. Keeping it in perspective, however, is crucial. Accordingly, the first part of my comment provides perspective on precisely where these data fit in. The second part addresses the specific results of the paper.

Some Perspective

This paper is about spot trading. It is important to keep this straight. For example, when the Bank for International Settlements writes of a \$1 trillion daily "foreign exchange market," many market segments are being lumped together: spot, forward, swaps, futures, and options (see BIS 1993). Care should be exercised when using aggregated BIS statistics to discuss the spot segment in particular. The authors themselves occasionally lapse (e.g., when discussing the market share of automated dealing systems, they refer to BIS data that are not from the spot segment alone).

Let me telescope further. Spot trading accounts for about half the foreign exchange total. Mark/dollar is the largest spot market by a margin, accounting for about a third of trading. Now, within spot markets, there are two main types of participants: dealers and customers. By *customers* I mean here any participant that does not provide two-way prices (e.g., corporate treasurers, investors, hedge-fund managers, liquidity traders, central banks, etc.). About 85 percent of spot mark/dollar trading is between dealers.¹

It is this interdealer trading that produces the D2000-2 data in this paper. Moreover, the data come from a particular type of interdealer trading, namely, brokered trading. There are two basic types, direct and brokered. Direct interdealer trades involve communication between the counterparties only. Price and quantity from these trades are not observed by others. In contrast, brokered trading involves prices that are advertised to dealers generally, as described in their section 4.2 (customers do not have access to interdealer brokers, electronic or otherwise). In spot mark/dollar, about two-thirds of interdealer trades are direct, and the remaining third are brokered.²

It is important, in my judgment, not to overemphasize the distinction between electronic (screen-based) trading and voice-based trading. More im-

Richard K. Lyons is associate professor in the Haas School of Business at the University of California, Berkeley, and a faculty research fellow of the National Bureau of Economic Research.

1. Table 1-B in BIS (1993) reports that customer-dealer trades account for about 12 percent of the total in spot mark/dollar. Calling the remaining 88 percent interdealer would be an overestimate, however. That report includes a third category of participant called *other financial institutions* that accounts for another 12 percent. This category includes nonreporting banks, which in many countries includes investment banks, some of which are important in dealing. (That dealers are included in this third category is evidenced by the significant brokered trading of this category; in general, only dealers have access to brokers.) Since this third category does include some "customers" by my definition (e.g., insurance companies and pension funds), 85 percent is a reasonable conjecture for the interdealer share.

2. Note that table VI of BIS (1993) does not report brokered shares for the spot market alone. For spot market data, individual central banks provide more information.

portant is the above distinction between brokered and direct trading, both of which have electronic and voice-based options. The D2000-2 system that the authors track competes with traditional voice-based brokerage.³ Another Reuters system, called Dealing 2000-1, is an electronic means for direct trading. Like voice-based direct trading, only the counterparties communicate when using Dealing 2000-1. Thus, speaking of "electronic dealing systems" without separating direct and brokered trading can be misleading since they involve very different dissemination of information.

With the above as background, here is a concentric rings model to organize the data sources referred to in the paper. There are three rings. The inner ring is direct interdealer trading. The Dealing 2000-1 data used in Lyons (1995; see also Lyons, chap. 5 in this volume) is from this inner ring. In mark/dollar, spreads in this inner ring are typically three to four pips for large banks when trading is active (London afternoon/New York morning). The second ring is brokered interdealer trading. The authors' D2000-2 data are from this second ring. Spreads in this ring are typically five to six pips when trading is active (here, I have large brokers in mind, which D2000-2 was not in 1993).⁴ The third ring is customer-dealer trading. Although transaction data from this ring are not currently available, my experience with dealers is that spreads are in the seven to twelve pip range for large customers (circa 1993). I view the indicative FAFX data as targeted at this third ring. That is, for most customers, this indicative series is the best real-time indicator of where the market is trading. Clearly, FAFX is not targeted at dealers since live broker quotes are more informative and they are easy to monitor continuously.

This leads to an issue that I do not believe has been addressed adequately in past work using FAFX: Exactly who inputs these indications at any given bank? And how? Stop to think for a moment about how rational it would be to pay a veteran dealer to input indicative quotes while trading, say, a billion dollars a day. No. It is much less expensive to hire a dependable young person to sit within earshot and intermittently type in a five or ten pip price based on where the dealer is actually trading. Better yet, why not build in some automaticity? For example, write a program that captures the dealer's firm Dealing 2000-1 quotes and widens them for customer consumption. A better understanding of this entry process would shed light on why the series has the properties it does. Let me suggest that the clustering of the FAFX spreads at

3. When I asked a spot mark/dollar dealer what market share broker systems like D2000-2 would have in five years, his response was essentially the following: "Currently, traditional brokers have about a third of the interdealer market, the rest being direct. Five years from now, I would guess that electronic brokers will have about half that third, traditional brokers the other half of that third, and direct will remain about two-thirds. There will always be a need for direct dealing." Admittedly, this is just one person's view, but tempered with experience nonetheless.

4. So why use a broker if direct prices are tighter (and brokers also charge a commission)? Smaller banks often do not have access to the tighter spreads among large banks. Large banks often prefer wider advertisement of their prices than bilateral direct quoting provides. Keep in mind that a large bank inputting the best bid, e.g., still buys at the bid side if a second bank hits that bid (it is the second bank that sells at the bid). Pretrade anonymity may also be valuable.

five and ten pips indicates that considerable automaticity is indeed built in (fig. 4.12).

Their Results

The authors present a fearsome array of results. In my judgment, this is appropriate given that their paper is the first study of its kind. This requires, however, that readers draw their own conclusions about what is most important to take away. The following is my take on it.

First, the paper brings to sharp relief the fact that there is no monolithic entity called *the spot market*. Even within mark/dollar, there are different ways to trade and different classes of participants. Consequently, there are many different sources of data. No one source provides a complete picture. The idea that D2000-2 is *the* market and therefore the ultimate benchmark is overwrought (and the authors are duly cautious here). That said, these are transactions data, and in that sense they represent market activity in a way indicative FFX data cannot.

With the authors' caution in mind, the three central take-aways appear to be the following. First, FFX provides an excellent image of the level of market price as it evolves over time. Second, FFX is a poor indicator of how market spreads vary over time. Third, FFX provides little information regarding trading volume (whether through entry frequency or otherwise).

Another result that I find interesting is their finding that the negative autocorrelation in returns disappears when transaction prices are used. The negative autocorrelation in FFX quotes is well documented and piqued enough interest that people had begun theorizing as to why it occurs. The fact that transaction prices do not exhibit this will surely affect how we think about it. Of course, as the authors point out, the reason that the quotes are autocorrelated while the transaction prices are not is an interesting topic in itself.

Two ways in which the paper might be clearer are the following. First, the text bounces a bit too much from comparative mode (D2000-2 vs. FFX) to focus mode (the properties of D2000-2 data per se). Although section 4.3 would appear to contain the comparative analysis, in fact the authors frequently compare the series elsewhere. This makes it difficult at times to know when the text is referring only to D2000-2. Second, in various places the text discusses "negative moving average" and "negative autocorrelation" without intending any distinction (as far as I can tell). Further, the term *reversal* is now commonly used in this literature to describe negative autocorrelation and would help readers less familiar with time-series work on returns.

References

- Bank for International Settlements (BIS). 1993. *Survey of foreign exchange market activity*. Basle: BIS, Monetary and Economic Department.
- Lyons, R. K. 1995. Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics* 39:321-51.