CHARLES GOODHART AND ANTONIS DEMOS*

Observations on the Swiss Franc/Dollar Spot Rate, via Quotations on the Reuters Screens

1. Introduction

By examining very high frequency data one is hoping to gain more insight into the microstructure of a market and consequently a better understanding of the lower frequency behaviour of the rate process. Furthermore, due to the large number of the ultra high frequency data, one can even after aggregating, still possess enough data in order to gain a better understanding of the returns distribution (see WASSERFALLEN/ZIMMERMANN (1985), WASSERFALLEN (1989) etc.). Moreover, intradaily frequency data provide an opportunity to address questions on the efficient processing of "news" (see eg GOODHART (1990), HARVEY/HUANG (1991)). By "news" we mean new information that is acquired by traders through the press and/or mass media and/or various relevant, or even irrelevant, sources.

It is common in the literature for variations in the arrival of "news" in financial markets to be measured directly from the data on the volatility of prices/returns (see eg ENGLE/NG (1991)). In one sense this approach assumes what needs to be tested, i.e. that "news" drives volatility. Moreover, the Autoregressive Conditional Heteroskedasticity (ARCH) effects commonly found in such financial series (see BOLLERSLEV/CHOU/KRONER (1992)), may well represent some combination of "news" arrival, i.e. the bunching of "news", and of "pure" volatility.

Here we describe some of the characteristics of the Swiss Fr/$ exchange rate. The basis for this exercise is a data set containing every single price entry on Reuters main forex screen from April 9th to June 29th 1989, a period of almost twelve weeks, eighty two days. During these days we collected every new "tick", or price entry, time-stamped, and showing the bank and the location of the branch making the entry (eg Citibank in Singapore). During this period there were some 148,079 price entries for the Swiss Franc/$ spot exchange rate included in our filtered data set, which, as will be reported in more detail shortly, involves on average some 2,500 price entries per week day.

The Swiss Franc/$ is one of the four most widely traded spot rates, alongside the Dm/$, Yen/$ and £/$, and is one of the four spot rates whose end hourly levels have been calculated for some years by Money Market Services, Ltd (see, for example, GOODHART/GIUGALE (1988)). A comparison of the frequency of price entries over Reuters screens in this period is shown in Table 1.

* Our thanks are due to Reuters PLC for advice and support, and particularly to Mark Jones and Ken Ferris of Reuters for explanations of Reuters systems. We thank Dr Russell Lloyd for assistance with software, and C V Bailey, D Barr and M D K W Foot for comments. Our special thanks are due to the editor and a referee of the Journal for substantial comments. They are not responsible for our interpretation of the data, nor for any residual errors which remain our own responsibility.

Table 1: Comparison of the Frequency of Price Entries over Reuters Screen

	Total	Average per Week Day	Average on Saturday	Average on Sunday
SwissFranc	148,079	2,496	2.5	8
Dm	305,589	5,120	7.5	74
Yen	165,530	2,775	6.0	70
£	135,741	2,283	4.5	28

Our purpose is twofold. First to provide detailed information on the frequency of new price entries (number of quotations within specific time interval) for this spot rate by time of day, by location and by bank. This will provide some evidence of when the market is most active, where it is most active, and who are the most active players on this particular network. Second, evidence is presented that activity, and consequently volatility, can be successfully approximated by the number of price revisions. On this basis, the implications of the recently developed informational models of ADMATI/PFLEIDERER (1988) and FOSTER/VISWANATHAN (1990) are discussed.

A short description of how the Reuters system works and the data set we analyse can be found in Appendix 1. Note that we were able to estimate the time of arrival of each new price quote to the second. We collected such data continuously (24 hours a day, 7 days a week) for almost 12 weeks, April 9-June 29th, with no breaks of any kind, until our computer blew a fuse on June 30th. The times are all at British Standard Time (BST). These are one hour behind Swiss time during the whole period.

There are many studies and exercises that can be done using these data (see eg GOODHART (1990), DASSIOS/DEMOS (1991)). Most of these involve examining how the actual bid/ask prices evolve over time; We and other colleagues are doing a series of such studies. But the purpose of this article is not to study the process of exchange rate price evolution, but primarily to explore the wider qualitative characteristics of the market, where, when and who. Thus the fact that the prices quoted are indicative [1], rather than firm, and that we do not have associated transactions prices and volumes is not in this case a problem.

What is more of a problem for our purposes in this part is that Reuters is only one of several purveyors of screen data, though one of the biggest (we do not have comparative data). Moreover, while the market is based on an electronic/telephone system, many participants do not choose to exhibit quotes over any of the screen systems. Furthermore much of the forex market goes through brokers, rather than via direct telephone dealing between two principal banks. Thus the BIS report on the survey by Central Banks in most industrialised countries and financial centres in April 1989, reported as follows [2]: "The role of brokers varies considerably from centre to centre. The share of foreign exchange business conducted through brokers ranged from a low of 18% of total net turnover in Bahrain to a high of 50% in Norway but averaged just under 40% in the three biggest market centres". Accordingly while we can use the data to show which banks are most active in using Reuters, that does not mean that locations and banks low on the list, or not (in the case of banks) on the list at all, are not active via other media of dealing.

Thus the relative share of foreign exchange brokers, and of other electronic screen systems, relative to Reuters may vary quite widely from centre to centre. In particular, as will be seen again later, there is some evidence the Reuters share is comparatively low in Japan, and perhaps to a lesser degree in the USA; whereas it is comparatively high in centres with a British connection, e.g. London, Hong Kong, Singapore, Sydney, Toronto, etc. Consequently the data on Where will also be subject to bias. Given the importance of Tokyo in the Asian markets, this suggests that the frequency of price revision may, perhaps, understate the true frequency in the Asian and American markets, relative to the European market.

Given the theoretical results on the hypothesis of mixture of distributions, the question of whether market activity affects volatility is of particular interest. Market activity is in general approximated by the volume of trade (see eg GALLANT/ROSSI/TAUCHEN (1990)). However, JONES/KAUL/LIPSON (1991) show that volume is a noisy and imperfect proxy for information arrival, and that the number of transactions is a better variable in a model with a fixed number of traders. However, there are no volume data available in the forex market (see eg GOODHART/DEMOS (1990)). Instead we shall use the frequency of quote arrivals over Reuters' screens as our proxy for market activity.

At this point we should repeat some pitfalls associated with the approximation of market activity by the number of quotations. For example, forex market participants report that at very busy/excited times the dealers may be too busy haggling over their phones to update their screen quotes. So, there may also be a tendency for the peaks (and perhaps the troughs?) of true underlying activity to be somewhat smoothed towards the average in these data series.

Moreover, in general the temporal pattern of the markets frequently differs from the temporal pattern of the "news" generation process. Markets often close almost entirely, eg weekends and the Tokyo lunch hour, or become very busy, while some "news" is continuously occurring. Although we would expect more "news" always to be associated with a higher frequency of quotes, as long as some markets are in operation, the functional form of this relationship eg linear, log-linear, etc., remains unknown.

The next question is whether it is permissible and appropriate to examine the contemporaneous interaction between quote arrival and volatility, or only to relate volatility to quote arrival using information available at t-1 and earlier. There are strong indications that the use of contemporaneous data on market activity virtually removes all persistence in the conditional variance in the returns series (see LAMOUREUX/LASTRAPES (1990) in daily stock returns, LAUX/NG (1991) in intra day currency future returns, and DEMOS/GOODHART (1992) in half-hourly forex spot returns in Deutsche Mark and Japanese Yen Dollar).

Our hypothesis is that, in this ultra-high frequency data set, the "causal" linkages will be found to be stronger from quote frequency to volatility when both are taken over the same short time interval, than vice versa. In GOODHART/DEMOS (1990), we argue that there are certain predictable temporal regularities in the Deutsch Mark foreign exchange market (eg the regular release of economic data at certain pre-announced times, the passage of the market through the time zones with market openings, lunch breaks (especially in Tokyo), etc.). These regularities are observed in the Swiss Fr/$ spot market as well (see below). Consequently temporal weekly daily and half-hourly dummies are added to all equations. As will be shown in Section 6, these two changes do make a difference to the results. Furthermore, the temporal dummies capture events (publicly announced news releases, market openings and closings) whose timing, though not generally their exact scale, is known in advance. Public news related to macroeconomic variables is simultaneously announced to all traders, at a time known in advance since the scheduled time of all economic related news is predetermined, and reported on another part of the Reuters system, the FXNB page. The stochastic element in such cases is the actual announcement, not the timing of it. In general, the majority of the US announcements are around 13:30 hours British Summer Time (BST), and the Swiss ones around 10:00 hours BST. Consequently, the relationship between the dummy variables and the characteristics of interest to us in the market, eg the bid-ask spread predominantly reflect the response of these variables to publicly known events. Per contra, the relationship between these variables, after conditioning on such temporal constants, will reflect private information to a somewhat greater extent.

Although the emphasis here is on the relationship between quote frequency and volatility, since it is a less-researched area, following DEMOS/GOODHART (1992) we examine the three-fold interrelationships between quote frequency, volatility and bid-ask spreads. The positive relationship between volatility and the spread is well-known in the large literature (see eg HO/STOLL (1983), BERKMAN (1991)). We suggested earlier that the absence of any significant ability of prior quote frequency to predict volatility implied that volatility may have incorporated both the contemporaneous evidence from quote arrivals and other sources of information: if so, we would not expect quote arrivals, either contemporaneous or lagged, to influence spreads, given volatility.

Subject to the above caveats , using the number of quotations within half hour, we shall explore the question of when the market is most active in the next Section, where the market is most active in Section 3, and who are the most active players in Section 4. In Section 5 we establish an empirical functional form between the number of quotations, as a measure of the market activity, volatility, and average spread and discuss the implications for the informational models of ADMATI/PFLEIDERER (1988) and FOSTER/VISWANATHAN (1990).

2. CHF/$: When?

During these 82 days during which we ran the computer, there were 148,079 entries. Naturally this was very much less over the weekends than on weekdays. However the new week really begins with inputs from Sydney and Melbourne in Australia, at about 22.30 hrs (BST) on Sunday night (their early Monday morning). Consequently the daily averages, slightly misleadingly, show more activity on Sunday, than on Saturday, but that is concentrated into the final hours of Sunday. Table 2 below gives the daily data, and the average value and standard deviation of figures for price entries on each week day. These figures are, however, affected by the incidence of holidays. Dates of holidays during our data period were collected from THE ECONOMIST ESSENTIAL DIARY INFORMATION (1989). While it might have been possible to test for the effect of each country/holiday date separately, (by the use of dummy variables), instead inspection of Table 2 and The Economist Diary suggest that there are four main holiday outliers in this data set, all in May, 1st, 4th, 19th and 29th. The average value and standard deviation of the individual weekdays, excluding such holiday outliers is also shown in Table 2.

Out of these outliers, due to holidays, only one of them, May 4th, was also a holiday in Switzerland. There were two other Swiss holidays in this period, April 17th (Zurich only) and May 15th, but both inspection of table 2, and subsequent regression analysis, revealed that neither had any apparent significant effect in reducing Swiss Franc forex deals. As can be seen from Table 2, weekend activity, as represented by the number of price revisions, (up till 22.30 hours on Sunday), is just about one thousandth of average weekday activity (2.5/2,500 = 0.1%).

Tuesday seems to be the busiest day, with some 200 extra price entries, and Friday, the least active day, with some 300 fewer price entries, than the average of Mondays, Wednesdays and Thursdays after excluding outliers. The standard deviations of the price entries on each weekday (again after excluding outliers) are all in the range of about 290/350 per day.

In order to examine the difference between average activity on each week-day in more quantitative form, we then ran a multiple regression over all week-days (excluding Saturdays and Sundays). Besides the weekday and holiday effects, there did appear to be an upwards trend in forex activity generally over these weeks (see GOODHART/DEMOS (1990) and (1991)), while the first week, and final three days of the sample, appeared unusually quiet. This may have been because April 1989 was a relatively quiet month in the forex market, while May and June were comparatively volatile, see BANK FOR INTERNATIONAL SETTLEMENTS (1990), Table 1. Volatility and activity are positively correlated. So we fitted dummies for each week-day, apart from Friday (the overall constant), a time trend, holiday dummies and dummies for the first week and final three days. The latter two dummies, at least, can be described as

Table 2: Daily Number of Entries

Sun	Mon	Tues	Wed	Thur	Fri	Sat
3	2358	2247	1924	2256	2172	3
16	2635	2399	2526	2193	2387	6
5	2119	2508	2248	2242	2217	2
14	726	2607	2432	1370	1896	1
3	2100	2524	2828	2496	2274	4
19	2315	2805	2531	3078	1536	6
7	2833	3244	2909	2930	2867	2
25	1740	3005	2922	3142	2731	0
0	2881	3028	3105	2278	2491	1
0	3291	3204	2680	2869	1920	0
1	2729	3044	2510	2916	2219	2
0	2814	2464	2499	2294		
Average per day including outliers						
7.7	2378.4	2756.6	2592.8	2505.3	2246.4	2.5
Standard deviations per day including outliers						
8.2	642.8	325.3	309.5	485. 5	361.9	2.0
Averages per day excluding outliers						
7.7	2607.5	2756.6	2592.8	2608.5	2317.4	2.5
Standard deviations per day excluding outliers						
8.2	359.9	325.3	309.5	359.7	297.7	2.0

First holiday Monday with 726 No. of Obs. Ist of May.
Second holiday Thursday 1370 No. of Obs. 4th of May.
Third holiday Monday 1740 No. of Obs. 29th of May.
Fourth holiday Friday 19th of May 1536 No. of Obs.

"data mining", (and we would certainly not want to extrapolate the time trend outside our short data period) but recall that our objective at this point is simply description (particularly of day-of-the-week effects), and not the construction and testing of causal hypotheses. The results were as shown in table 3, with robust to heteroskedasticity t values shown below.

The results confirm Table 2, as indeed Tuesday is the busiest of the week-days. However, although this is in agreement with the other three major currencies (see GOODHART/DEMOS (1991)), it is quite surprising since the majority of the US economic related announcements took place during this period primarily on Fridays [3]. As it is widely believed that bilateral to Dollar currencies are mainly US news driven markets (see ITO/ROLEY (1987), ENGLE/ITO/LIN (1990) on Yen/$ rate , GOODHART (1990) on Dm/$, £/$ and Yen/$), this discrepancy between the frequency of entries and activity might, at first sight, indicate that activity in the markets is rather poorly approximated by the number of entries (but see section 5).

Table 3: Results of multiple regression over all week-days

Constant	Week Days			
Fri	Mon	Tues	Wed	Thurs
1948	302	497	323	315
(17.5)	(2.4)	(4.1)	(2.7)	(2.6)
Time	Holidays			
Trend	May 1	May 4	May 19	May 29
12.3	-1721	-1127	-793	-957
(5.3)	(-5.9)	(-3.9)	(-2.7)	(-3.5)
First Week	Final 3 Days	$\bar{R}^2$=0.67		
-115	-602	DW = 1.62		
(-0.76)	(-3.3)	F(11,47) = 11.59		

Such daily averages represent a comparatively aggregated, broad-brush picture of the time-varying frequency of price entries. To investigate in more depth we have, therefore, chopped up each day into half-hourly periods and we reproduce the histograms of the average frequency of price revision in each half-hour in each week-day in Charts 1-5. There are a number of apparent features in these Charts. First, the overall time pattern on each week day is remarkably similar, with three main peaks connected with the opening of the three main market areas (Asia/Tokyo at 01.30, Europe/London at 07.30/08.00, and N.American/New York at 14.00). There is a long (Pacific) trough from about 17.30 (shortly after the close in Europe) till 00.00 with the

Asian opening, and a dramatic collapse in dealing from 04.00 till 05.30 (in the Tokyo lunch hour, of which more later).

The lower average frequency of trading on Friday in Swiss Franc seems to be compounded of two factors, first a generally lower level of European trading (whereas trading in the Asian time zone is, if anything, higher on Fridays than on Wednesdays/ Thursdays), and, second, a quicker and more marked decline in trading activity after the New York and European trading peak around 14:00-14:30 BST. The Anglo-Saxons may, perhaps, show a greater allegiance to the POETS (Push Off Early, Tomorrow's Saturday), than the Asians.

There is, however, an alternative explanation for the reduction in trading activity on Friday afternoons. Foreign exchange dealers are reluctant to run open positions across a weekend because they would then be exposed to two days' political developments without being able to cut out a deteriorating position. So Friday is the day when dealers largely square off their positions, and once they are square they try to withdraw from the market so as to avoid being given a fresh exposure. So as Friday progresses trading becomes increasingly thin and it becomes difficult for laggards to trade out of their positions. The problem is least acute for Far East

Charts 1-5: Average Frequency of Price Revision in each Half-Hour in each Week-Day

CHART 1
MONDAYS CHF

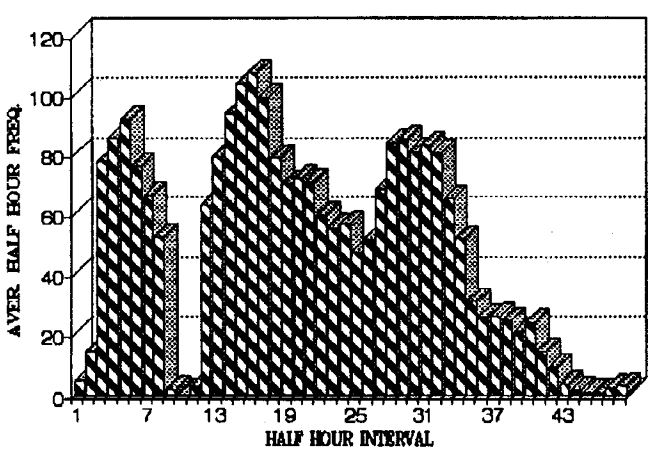

CHART 2
TUESDAYS CHF

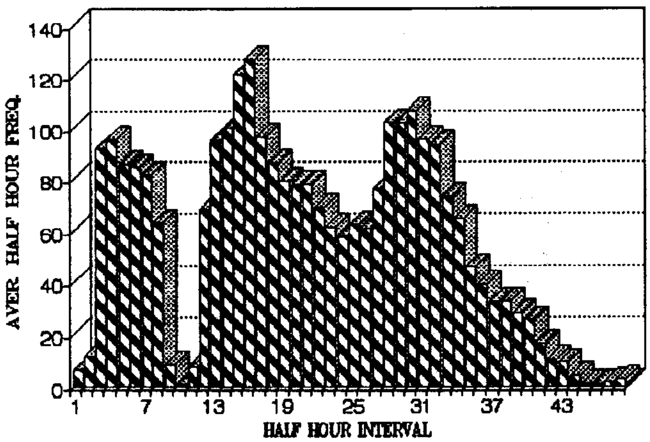

CHART 3
WEDNESDAYS CHF

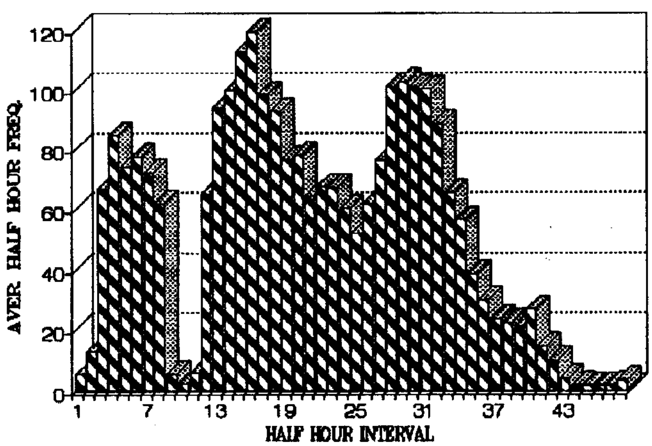

CHART 4
THURSDAYS CHF

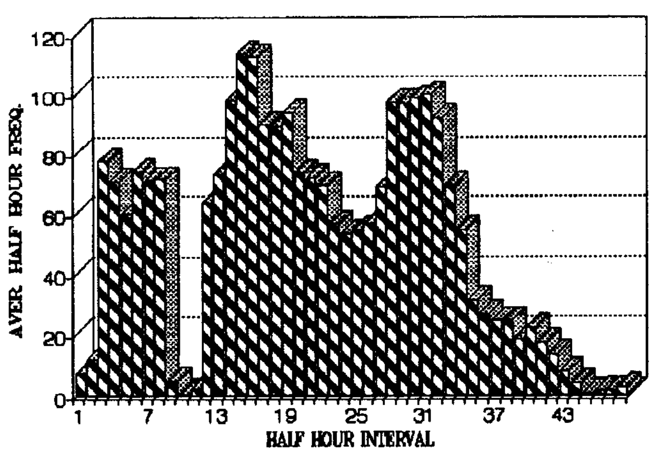

CHART 5
FRIDAYS CHF

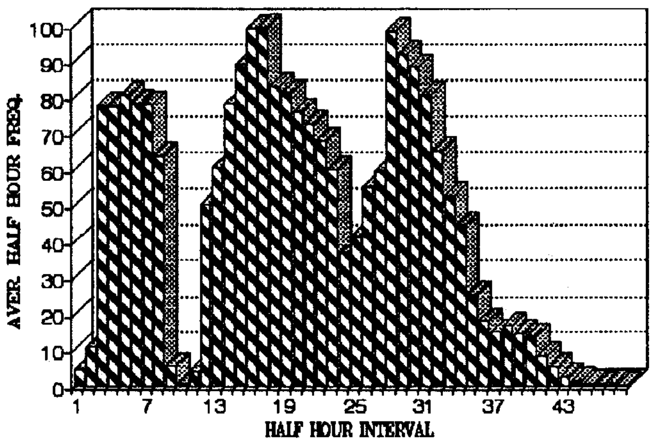

traders, for if they have undesirable exposures left at the end of their Friday they can simply trade into their evening, i.e. into the active European morning, or leave orders with their branches and subsidiaries in Europe. However, that course is not open to New York banks, where there is little activity in their afternoon in any case, and particularly so on a Friday afternoon. Consequently they are forced to square up early in their day and stop trading. So the reduction in trading late on Fridays will, at least in part, be a reflection that in fact there is one, global market and not separate geographical markets, and that by accident of time zones the Anglo Saxons are trading towards the end of the global Friday.

There are also two further reasons for the apparent discrepancy between activity driven by US announcements and frequency of entries. Firstly, some of the Friday announcements took place when only the American markets were effectively in operation i.e. after 16:00 hours BST, whereas most of the Tuesday announcements took place when the European markets were in operation as well. Hence the effect of the US late announcements should be apparent in the opening of the Antipodean and/or Asian markets on Mondays. However, the forex market is fast in digesting economic news. A piece of economic news has an affect, if at all, in the following half an hour at the most (see GOODHART (1990)).

Secondly, a careful inspection of Chart 5 reveals that although Friday as a whole is a quiet day, with a morning peak which is at least 20% lower than Tuesdays' morning peak, Friday early afternoon peak is as high as the morning one and more importantly as high as Tuesdays' afternoon peak. This could be very well due to the US economic announcements. Finally, as we already mentioned there may be a certain downwards bias of the frequency entries, as a measure of activity, at the peaks, for during busy periods traders could be preoccupied with making deals rather than updating their screens.

Saturdays [4] show evidence of some small continuing forex trading in Asia and rather more frequent trading in the European time zone (mainly from the Mid-East as will be demonstrated later). The sanctity of the North American weekend is clearly apparent. On Sundays there is only a smattering of Mid-Eastern quotes, before the Antipodes reopen at 22.30pm. Due to the small number of entries during the weekends (see Table 2) and the lack of any economic related announcements it is very difficult to undertake any analysis in terms of frequency of entries and activity and/or volatility. However, "news", as defined here, occurs during this period of the week as well. The effect of that should be apparent at the Monday opening of the various markets. According to informational models of ADMATI/PFLEIDERER (1988) and FOSTER/VISWANATHAN (1990) one should expect to observe an increase in volatility during the opening times of the markets and these increase should be more apparent in the Monday openings (see also HARVEY/HUANG (1991)). As now activity and volatility are positively correlated, one should expect an increase in the number of entries at the openings and a gradual decline for the rest of the morning. This should be more apparent for Mondays than for any other of the week-days. A careful inspection of Chart 1 reveals that this is the case for the Asian market (see below for more on this topic).

As a further measure of the differences/similarity of activity patterns, intra-daily, between the various week-days, we have taken the percentage of each day's total price revisions occurring in each half hour of each week-day, and then examined the resulting cross-correlations between days of the week, excluding Sundays, for which there were too few quotes (before 22.30pm) to make a useful comparison. The results are shown in Table 4.

Table 4: Cross-Correlations between Days of the Week

	MOND	THUES	WED	THUR	FRI	SAT
MOND	1.000					
TUES	0.989	1.000				
WED	0.984	0.988	1.000			
THUR	0.971	0.981	0.985	1.000		
FRI	0.972	0.968	0.969	0.969	1.000	
SAT	0.248	0.230	0.219	0.262	0.234	1.000

The correlation between half-hourly % frequency of price revisions is close for all week-days, in all cases over +.965, with Friday being very slightly less similar, probably due to weekend effect (see above). The correlation between the half-hourly intradaily pattern on the weekdays and on Saturdays is, as would be expected, far less close.

These cross-correlations for the Swiss Franc were slightly higher than for the Yen or £, but slightly lower than the Dm.

We then examined the correlation on each day between the half-hourly % frequency of the Swiss Franc and of the other three currencies. The results were as shown in Table 5.

Note the extremely high weekday correlation between the daily pattern each day of the Swiss Franc and Dm and £, and rather low correlation with the intradaily trading pattern of the Yen. This could be, very well, attributed to the common factor that these rates have i.e. news related to the US$, and that the three currencies with the high correlation are European. This result, however, is reversed over the week-end, Saturday and Sunday, (perhaps because the Dm and £ are more commonly traded over the week-end in the Middle East centres than the Yen and CHF?).

3. CHF/$: Where?

We turn next to examine the location (where) of bank branches inputting CHF prices into Reuters. In Table 2 we present the listing of entries by the financial centre from which each CHF entry was made.

It is most important to repeat the caveats about the volume of Reuters screen entries not providing an accurate representation of total activity [5]. There are various biasses, i.e. the varying share of Reuters in total screen activity, and the varying relationship between screen entries and actual total transactions. With these caveats firmly in mind, one may inspect Table 6.

Some 25 centres quoted prices for CHF/$, fewer than for Dm/$ (38), but more than for Yen/$ (23) or £/$ (21). This was partly because several Swiss centres (Basle, Geneva, Lugano) put in price entries for CHF/$, whereas only Geneva had quoted Dm/$, and none of these three for Yen/$ or £/$. Zurich was by far the largest centre for CHF/$ entries, having 60% more entries than the next largest centre, London. The four Swiss centres together, (though Lugano had few entries), made 28% of all CHF/$ price entries, whereas London had a larger propor-

Table 5: Correlation on each Day between the Swiss Franc and the Other Three Currencies

	SUN	MON	TUES	WED	THUR	FRI	SAT
CHF and DM	0.951	0.974	0.955	0.976	0.973	0.973	0.680
CHF and YEN	0.976	0.847	0.822	0.846	0.824	0.831	0.847
CHF and £	0.868	0.954	0.961	0.985	0.972	0.964	0.508

Table 6: Numbers of entries per location

Location	No. of Quotations
ZUR	33086
LDN	20703
SIN	20608
H.K	20594
N.Y	18962
TOK	7750
BAS	5842
OSL	5841
TOR	3494
GVA	2711
COP	1596
BAH	1448
CGO	1435
SYD	1005
BRU	845
KUW	558
MEL	528
FFT	412
VIE	319
LUG	99
S.F	84
L.AV	79
BOS	54
JED	19
DUB	7

tion of £/$ entries (34%) and the German centres and Tokyo had a smaller share of Dm/$ and Yen/$ business (18% in both cases).

After Zurich the next four (of the big five centres) London, Singapore, Hong Kong and New York all had roughly the same volume of CHF/$ entries, at about 20,000, with Tokyo trailing quite a long way behind with 7,750. This latter is again almost certainly due to the underrepresentation of Reuters in Tokyo. Oslo and Copenhagen make comparatively a lot of price entries in CHF/$, and Toronto and Sydney maintain their position as middle sized currency-trading centres (see GOODHART/DEMOS (1990)). The Middle East centres, with the exception in this case of Abu Dhabi, (i.e. Bahrain, Kuwait, Jeddah and Dubai) play their usual small, (partly week-end), role. In the USA Chicago and Boston play a comparatively larger role with 1,435:54 entries, than with Dm (197:0), £ (387:0) or Yen (107:0), whereas the West Coast centres, Los Angeles and San Francisco play quite a small role (79:84). In the Antipodes Sydney and Melbourne are the only centres active in CHF/$ (1,005:528), and Wellington does not participate in making price quotations over Reuters for CHF/$. Brussels, Frankfurt and Vienna are the other centres in Europe entering Reuters quotes for CHF/$.

We can, perhaps, obtain some indication of the various biasses by comparing our results in Table 6 with the data for the Swiss Franc composition of total forex currency business, presented in Tables C-2 and C-3 of BANK FOR INTERNATIONAL SETTLEMENTS (1990), obtained from the Central Bank survey in April 1989, in which survey the Deutsche Bundesbank and Institut Monétaire Luxembourgeois did not fully participate, whereas data for the USA is not available.

Subject to the limited amount of information presented in this survey the general conclusion is that the ordering in Table 6 is not very different of the ordering in terms of volume of transactions. The only really noticeable differences are, first the complete absence of French centres from Table 6, as compered with the 7th position that France occupies in BIS survey. Second, although in Table 6 Singapore and Hong Kong have more or less the same number of entries, BIS survey reveals that the Hong Kong transactions' volume is around 70% of the Singapore volume.

As the world spins around each day, various centres in each part of the world become active and then go to sleep at differing times. We illustrate this by a histogram of frequency of entry for each half-hour, (averaging over all days in the work-week), for the six centres with the largest number of Reuters' quotes, viz Zurich, London, N.Y., Hong Kong, Singapore and Tokyo for Swiss Franc. These are shown in Charts 6-11.

These Charts show a differing pattern of temporal activity in the three main market areas (Asia, Europe and N. America).

Incidentally the caveats about likely biasses in Table 2 would seem much less troublesome in the interpretation of these charts. The three Asian markets

Charts 6-11: Frequency of Entry for Each Half Hour for Different Centers

CHART 6
CHF ZURICH

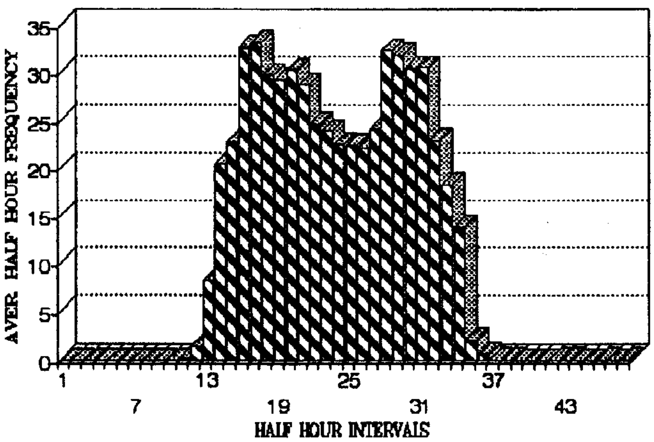

CHART 7
CHF LONDON

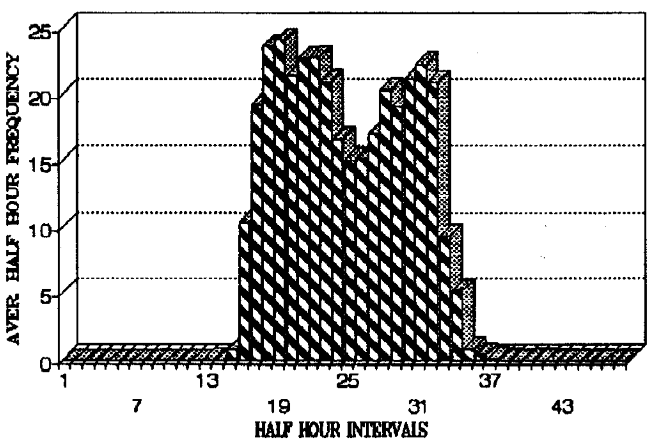

CHART 8
CHF SINGAPORE

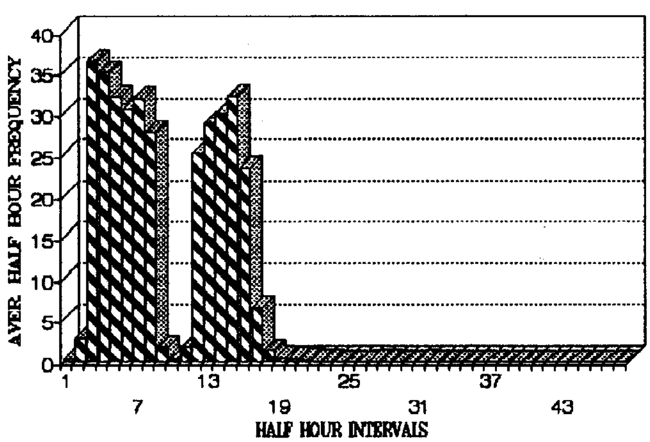

CHART 9
CHF HONG KONG

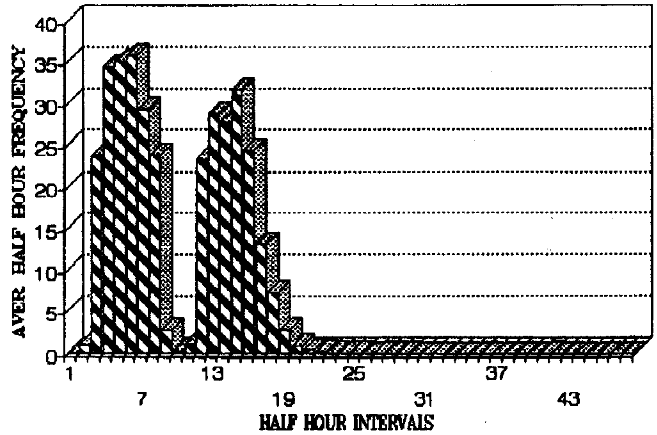

CHART 10
CHF NEW YORK

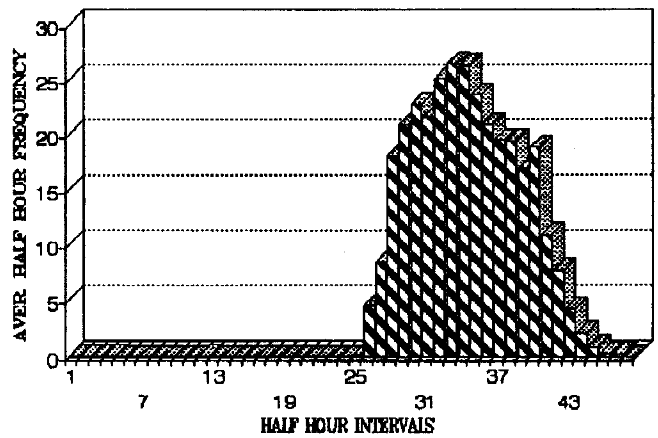

(Tokyo, Singapore and Hong Kong Charts 11, 8, and 9 respectively) have remarkably similar time profiles, although HK is one, and Singapore two, hours behind Tokyo. The most obvious feature is the collapse of trading during the Tokyo lunch hour (04.00-05.30). Since there is an equivalent trading halt at the same (GMT) time in Australia, Singapore and HK, despite differing local times, the evidence would seem to point towards Tokyo hegemony in forex activity in this area. In so far as the Tokyo market dwarfs its neighbours in volume terms, with forex traders generally preferring to trade when the market is active and to have lunch when it is quiet, so traders in other Asian markets will operate by Tokyo's clock rather than their own. The failure of our Reuters price revision frequency data to show any such Tokyo domination may reflect Reuters' market weakness in Japan.

Apart from the lunch hour, trading is bi-modal, with, perhaps, less concentration at the opening (and closing) than current financial market theories (ADMATI/PFLEIDERER (1988), FOSTER/VISWANATHAN (1990)) would have suggested. The obvious exception is Singapore, where the daily pattern would almost completely agree with the private informational models. Trading is somewhat busier in the Asian morning before lunch than afterwards, probably due to the US economic announcements of the previous day. Trading then declines quite rapidly in Tokyo after 07.30, after 08.00 in Singapore, and after 08.30 in Hong Kong.

In Europe (Chart 6, Zurich: Chart 7 London), trading is again bi-modal, but the mid-day slowdown, 12.00-13,30 (BST), is much less marked. Even ignoring the difference in lunch habits, the London working day 07.30-17.00 is longer than in the Asian market (01.00-07.30 Tokyo; 01.30-08.00 Singapore; 01.30-08.30 Hong Kong). Again the morning peak is above that in the afternoon, but not by much. The afternoon peak comes late, perhaps coinciding with the build-up of activity in New York.

Unlike the Asians and Europeans, New York (Chart 10) forex dealers do not appear to take a lunch-break from their consoles. Alternatively, as one commentator suggested, "The market perception is that New Yorkers go to lunch and do not come back. This is confirmed by ..... the steady decline in US activity after 17.30 BST (12.30 EST) and the absence of an afternoon peak in New York, and arises from the fact that, after London's close, the reduced size of the market makes New York trading relatively thin". Trading activity appears clearly unimodal with a steady build-up from an early opening at 13.00 (BST) to a peak at 17.30 (BST), followed by a steady deadline in activity. This is a clear indication against the privately informed trader theories but is in favour of the public announcements model.

As the forex market is almost continuously in operation, it is very difficult to distinguish between public and private information models (but see HARVEY/HUANG (1991)). The only exception is the Pacific through, when effectively no markets are in operation. Consequently, the behaviour of the Asian markets could offer the chance for distinguishing between the two theories. However, as we already discussed there is no clear cut answer to the question.

Besides looking at the six main centres by themselves, we have also looked at the entries, by half-hour, from all locations grouped into six main market areas, as follows:

(1)	Antipodes:Sydney, Melbourne	(2)
(2)	Asia: Hong Kong, Tokyo, Singapore	(3)
(3)	Middle East: Bahrain, Kuwait, Dubai, Jeddah	(4)
(4)	Europe: Zurich, London, Frankfurt, Copenhagen, Basle, Geneva, Lugano, Brussels, Oslo, Vienna	(10)
(5)	N. America (E): New York, Toronto, Boston, Chicago	(4)
(6)	N. America (W): Los Angeles, San Francisco	(2)

Histograms showing the average weekday half-hourly frequency of entries from these areas onto Reuters are not shown for space economy. These mostly confirm the results described in Charts 6-11, and are also virtually identical to the patterns observed for the Dm, £ and Yen and reported in our earlier papers. Noticeable is, for example, the slowdown in Antipodes trading during the Tokyo lunch hour, (well after their own). Asia and Europe have strongly bimodal frequencies, with little difference in activity before and after lunch, whereas N. America (E), Antipodes and the Mid-East are uni-modal. Only in the case of the Mid-East, however, does the peak come early in the trading day. The number of observations from N. America (West) was very small. Even so, there is evidence of an attempt to maintain both an early link with trading on the East Coast, and with rather more activity at a time of a late link with the Antipodes from 21.00-24.00 hours (GMT).

We also looked at price entry patterns over the weekend. Since such activity was so, comparatively, minuscule, we divided the 48 hours, 00.00

Friday/Saturday till 00.00 Sunday/Monday into eight equal six hourly periods, and then calculated the histograms of the frequency of price entry, by the six main market areas, in these periods. The overlap from Friday action in North America into early Saturday, and the entry of the Antipodes late on Sunday, is clear. Apart from this, North America totally abandons the office on weekends, unlike Asia which still has a smattering of trades on Saturday, and the Mid-East, which keeps going on both Saturday and Sunday. Indeed, comparing the appropriate Charts (and after due allowance for the difference of scale, ½ hour vs 6 hour), the extent of Mid-East activity is much the same on the week-end as on the average weekday (excl. Friday).

To summarize and conclude this Section:

(1) Cultural differences in lunch-hour behaviour show up strikingly. In Asian market space this is dominated by the Tokyo lunch hour, presumably because Tokyo is the dominant market, even though that does not show up either in our price frequency series or BIS survey.

(2) There appears to be much less bunching of screen based activity at market openings than some current theories might have predicted. Asia and Europe are bimodal, with the pace of afternoon activity only slightly lower than in the morning. The US, Antipodes and Mid-East are all uni-modal, but only the latter has a peak early in the trading day. Unlike Stock Markets, however, there is no formal opening, or closing, in the forex market, which is continuous. Most of the theory on the bunching of activity (at market openings) has been based on discrete Stock Markets, not on continuous forex markets. Indeed on Mondays, when the market does have a discrete opening time, activity is not as lethargic, (relative to Tuesday - Thursday) early in the day (as it becomes after about 12.00 BST).

4. CHF/$: Who?

We turned next to an examination of which banks made most use of Reuters screens for in-putting quotes for the CHF/$. It is essential to remember that a large proportion of banks may prefer not to input quotes on any screen, or will use screens other than Reuters. The lists below must not be regarded as a comparative measure of activity in the forex market. The data are exactly what they purport to be, a listing of frequency of entry, over our data period, of quotes for the spot CHF/$ by banks using Reuters.

Banks use the Reuters screen system for two main purposes. One is to keep their corporate clients posted with an indication of the latest rates, and the other is to advertise themselves to the market (corporate and inter-bank) by keeping their name on the FXFX screen. The extent to which banks pursue either strategy may vary considerably, and the variations may have nothing to do with the importance of the bank as a dealer in foreign exchange. Indeed, a bank's use of Reuters may reflect temporary, human elements such as having a keen trainee dealer who has the responsibility for updating the rates.

With these caveats firmly in place, we may turn to Table 7, which shows the twenty five branches with the largest number of individual entries for each of these two currencies. There were some 116 branches round the world which input such entries for the CHF/£ during our time period. This is around 40% of the number of participants reported in WASSERFALLEN/ZIMMERMANN (1985), and WASSERFALLEN (1989). The "market" as a whole was moderately concentrated, the largest entrant making 5.7% of the entries, the top five 25.1%, the top ten 44.4%. The Herfindahl index stands at 0.0292, implying that, if all the branches undertook equal price entry activity, then it would take 34 of them to make an equivalent number of entries. This is a considerably higher degree of concentration than for Dm, though only slightly larger than for £/$. This can also reconfirm earlier claims that UBS in Zurich is one of the most active participants of the CHF/$ market (see WASSERFALLEN/ZIMMERMANN (1985), and WASSERFALLEN (1989)).

Simple inspection of Table 7 reveals that many international banks input quotes from several branches around the world. We can, of course, aggregate over all branches to show which banks were most active in this particular respect. This is shown in Table 8 for the 25 banks with the largest number of CHF/$ entries. There were 70 banks in all making such inputs for CHF again a smaller number than for Dm (135), Yen (110) or £ (90). Concentration was naturally somewhat greater with the top 1, 5, and 10 banks making 16.4, 59.5, 77.4% of the entries respectively. The Herfindahl index stands at 0.0842.

As in the case of our earlier studies, (on Dm/$, Yen/$ and £/$), Hong Kong and Singapore are more overrepresented on the Reuters data, in terms of volume of activity, than in the Central Bank survey. Once again there was no Reuters entry at all from Paris for the Swiss Franc/$ exchange rate, but the Central bank survey indicates that Paris is a (minor) centre for spot rate dealings in CHF/$. On the other hand the Reuters' data show a considerable volume of price entries, for CHF/$, from Oslo, Copenhagen and Bahrain, (and a smaller volume of such entries from Brussels), which find no reflection in the Central Bank survey.

Overall the correspondence between the two data sets is moderate, though (apart from the complete omissions from the CB survey), most independent observers would probably attribute most of the differences between the two data sets to various biasses, already noted, in the Reuters sample.

Not surprisingly the Swiss banks are the most active in the CHF/$ spot exchange market, with two banks (Swiss Bank and UBS) taking the top two places, and Credit Suisse in the fifth place. Two big US banks, Morgan and Citibank, are third and fourth in this list, with Chemical sixth. Norske bank also has a large presence (7th).

Finally, for this section, we returned to the three bank branches with the largest number of individual entries, and charted their average daily number of

Table 7: Number of entries for individual branches

Name of Branch	Location	Number of Entries	Proportion of all Entries
C SUISSE	ZUR	8529	0.057598
U B S	ZUR	7921	0.053492
CHEMICAL	LDN	7209	0.048683
MORGAN	N.Y	7005	0.047306
MORGAN	SIN	6519	0.044024
CITIBANK	SIN	6511	0.043970
SWISS BK	BAS	5842	0.039452
NORSKE	OSL	5841	0.035285
SWISS BK	TOK	5225	0.035285
SWISS BK	ZUR	5085	0.034340
CITIBANK	N.Y	4654	0.031429
SWISS BK	H.K	4212	0.028444
ZUERCHER	ZUR	3884	0.026229
MORGAN	H.K	3855	0.026033
VOLKSBK	ZUR	3472	0.023447
DEUTSCHE	N.Y	3072	0.020746
BARCLAYS	LDN	2836	0.019152
U B S	LDN	2631	0.017768
CHASE	H.K	2488	0.016802
NATWEST	LDN	2241	0.015134
CHEMICAL	N.Y	2179	0.014715
ROY SCOT	LDN	2041	0.013783
ROYAL BK	TOR	2025	0.013675

Table 8: Number of entries per bank

Bank Entries	Price Entries	Proportion of Total
SWISS BK	24330	0.164304
U B S	18229	0.123103
MORGAN	17735	0.119767
CITIBANK	14521	0.098063
C SUISSE	13366	0.090263
CHEMICAL	9388	0.063399
NORSKE	5841	0.039445
VOLKSBK	3932	0.026553
ZUERCHER	3884	0.026229
CHASE	3335	0.022522
DEUTSCHE	3072	0.020746
BARCLAYS	2836	0.019152
NATWEST	2304	0.015559
ROY SCOT	2041	0.013783
ROYAL BK	2025	0.013675
CHICAGO	1907	0.012878
A B N	1798	0.012142
ANDELSBK	1596	0.010778
HARRIS	1435	0.009691
B C I	1366	0.009225
MAN HAN	1332	0.008995
DRESDNER	1141	0.007705
BK TOKYO	1123	0.007584
BK LEU	853	0.005760
DAI-ICHI	823	0.005558

entries, to see if these all moved quite closely in line with the aggregate pattern shown on Table 3 (n.b. no attempt was, however, made to correct for individual country holidays, though that could be done). The results for our three most active branches are shown in Charts 12-14.

In this case, unlike the other currencies that we have examined, the individual bank branches have weekly patterns not far from that of the aggregate. Indeed Chemical (London) is closely similar except that the aggregate dips more on Mondays. Credit Suisse (Zurich) is a bit unusual in that its volume was lower in mid-week than on Mondays/Fridays, the exact opposite of the weekly pattern of the UBS (Zurich). Even so, there was less evidence of idiosyncratic bank behaviour in these instances than in similar exercises for Yen/$ and £/$. The Friday unusually high activity of these three branches, as compared with the aggregate weekly pattern, is mainly due to that these branches are based in Europe and consequently are not subject to the coming weekend effect (see above).

Charts 12-14: Average Daily Number of Entries for the Branches with the Largest Number of Individual Entries

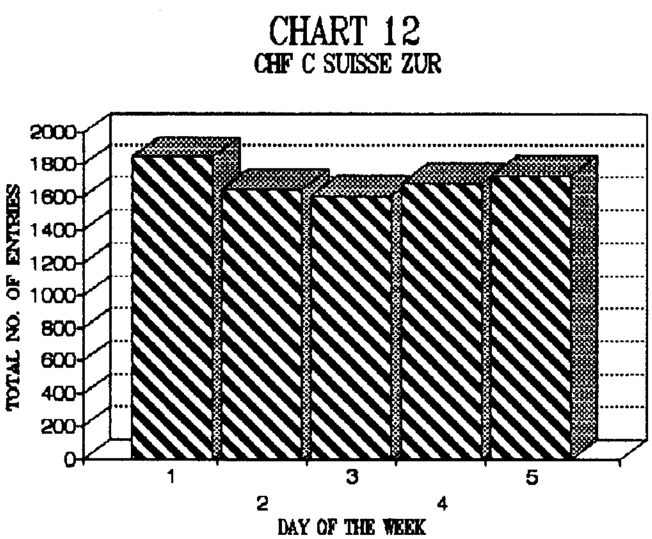

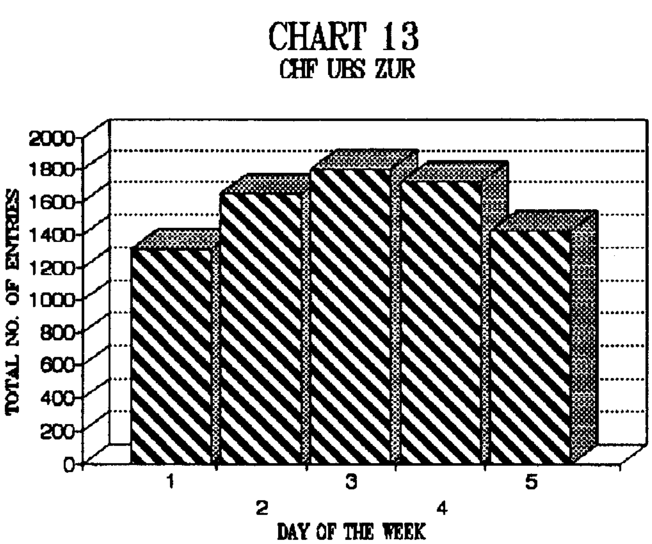

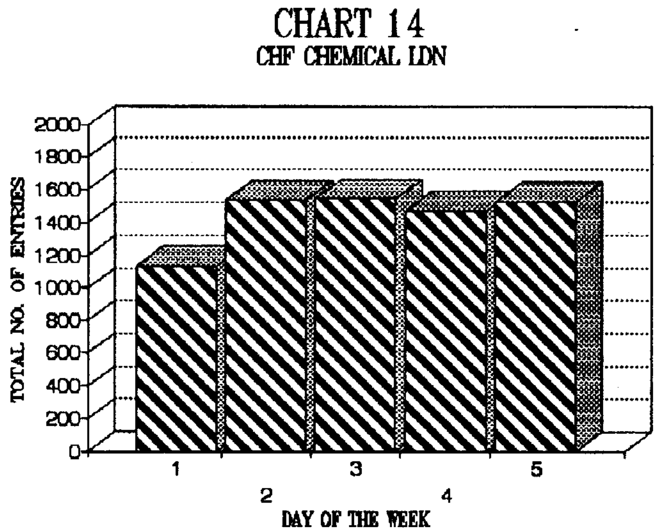

5. The Relationship between Quotes, Volatility and Spread

Following DEMOS/GOODHART (1992) the number of quotations, the average spread, and the standard deviation of the percentage difference of the rates quoted within the 48 half hour intervals of the week-day are recorded. This resulted in 2500 half hourly observations. In fact, 5 out of the 12 weeks were chosen for analysis, avoiding weeks with public holidays in the USA or Switzerland, and the following simultaneous equation system is estimated:

$$s_t = Dum. + a_{12} sp_t + a_{13} n_t + a_{14} s_{t-1} + a_{15} s_{t-2} \qquad (1)$$

$$sp_t = Dum. + a_{21}s_t + a_{23}n_t + a_{24}sp_{t-1} + a_{25}sp_{t-2} + a_{26}sp_{t-3} \quad (2)$$

$$n_t = Dum. + a_{31}s_t + a_{32}sp_t + a_{34}n_{t-1} + a_{35}n_{t-2} \quad (3)$$

where s_t, sp_t and n_t are the standard deviation of the percentage change of the exchange rate, the average spread, and the number of quotations within the t^{th} half hour interval, and Dum. are weekly, daily, and half-hourly dummy variables.

Table 9: Log-likelihood Values of the Reduced Form System

Value of Equation	Stand. Deviat. Equation	Aver. Spread Equation	No. of Quot. Euquation
1.0	-1714.08	-1884.23	-4984.90
0.6	-1476.28	-1831.25	-4607.06
0.5	-1438.10	-1824.69	-4542.13
0.4	-1408.10	-1820.67	-4498.71
0.3	-1386.28	**-1819.09**	**-4480.86**
0.2	-1372.42	-1819.88	-4496.04
0.1	**-1366.54**	-1822.94	-4550.36
0.0	-1368.79	-1828.15	-4647.39

Bold characters indicate optimum values

However, there is no apparent reason why the functional form between these variables should be linear, rather than, say, log-linear. Consequently, we left the data to decide on this, by transforming the three variables, using the Box-Cox transformation, and maximising the log-likelihood function of the reduced form of the system above with respect to the Box-Cox exponent. Notice that the linear and log-linear functional forms are special cases of this transformation.

From Table 9 it is apparent that neither the linear relationship, exponent equal to one, nor the log-linear, exponent equal to zero, is the optimal functional form. For the optimal values, 0.1 for the standard deviation, and 0.3 for the average spread and number of quotations, we performed some diagnostic tests, reported in Table 10. In particular, the WU (1973) and HAUSMAN (1978) F test for the exogeneity of the three variables, with one exception, is rejected. Consequently, the following

Table 10: Diagnostic Tests for the Unrestricted VAR

Test	Stand. Dev. Equation	Aver. Spread Equation	No. of Quot. Equation
Wu-Hausman	8.33	3.00	17.57
Basmann	5.562	2.41	1.165
Omitted Var. (2 More)	1.16	1.97	2.09
Serial Cor. (Order 10)	1.53	2.83	1.33
Normality	677.1	2136.5	386.7
ARCH(5)	55.33	102.4	61.01

simultaneous equation system is estimated, by two stage least squares. The estimates of the structural parameters and their heteroskedastic robust standard errors are presented in Table 11.

$$s^*_t = Dum. + a_{12}sp^*_t + a_{13}n^*_t + a_{14}s^*_{t-1} + a_{15}s^*_{t-2} \quad (4)$$

$$sp^*_t = Dum. + a_{21}s^*_t + a_{23}n^*_t + a_{24}sp^*_{t-1} + a_{25}sp^*_{t-2} + a_{26}sp^*_{t-3} \quad (5)$$

$$n^*_t = Dum. + a_{31}s^*_t + a_{32}sp^*_t + a_{34}n^*_{t-1} + a_{35}n^*_{t-2} \quad (6)$$

where $s^*_t=(s^{0.1}_t-1)/0.1$, $sp^{*t}=(sp^{0.3}_t-1)/0.3$, and $n^*_t=(n^{0.3}_t-1)/0.3$, and again *Dum.* are weekly, daily, and half-hourly dummy variables.

In Table 11 the estimates of the above system are displayed. Notice that the average spread in the volatility equation has significant explanatory power, and the same applies for the standard deviation in the average spread equation. However, the number of quotations does not have any significant direct effect on volatility and only marginal effect on the average spread. This could imply that there is some "causation" from the number of quotations to average spread. In fact, this is not completely true since there is a strong correlation between the residuals of the three equations. Indeed the correlation between the residuals of the first and the third equation is 0.30 and between the second and third is 0.17.

Notice, also, that in the volatility equation the coefficients of the first and second order lagged standard deviation are 0.19, and 0.10 respectively,

Table 11: Estimated Coefficients (a_{ij}) for the imtimated Simoultaneous Equation System in Equations (4)-(6)

i\j	1	2	3	4	5	6
1		0.421	0.019	0.190	0.105	
		(4.89)	(0.93)	(4.32)	(2.90)	
2	0.247		-0.014	0.416	0.077	0.067
	(2.73)		(-1.67)	(6.52)	(1.60)	(1.72)
3	-0.142	-0.086		0.416	0.060	
	(-0.33)	(-0.25)		(9.91)	(1.89)	

implying a mild autoregressive heteroskedasticity effect. However, this is not the case when average spread and number of quotations are excluded from this equation. In such a case the OLS estimates are 0.29 and 0.15 respectively. Moreover, the addition of our dummy variables further reduces the ARCH type effect in the series. If the above simultaneous equation system is estimated without the dummy variables, the coefficients of the first and second order lagged volatility are now 0.30, and 0.16 respectively.

Reverting to the system where half-hourly, daily and weekly dummies are included notice that the constant represents the last half hour of the last Friday in the sample. During this half hour markets are closed and only a few traders, if any at all, input quotations. Therefore, the constant in the estimation reflects, on average, the smallest number of observations in the sample, but not necessarily the lowest level of volatility or the smallest average spread. Let us concentrate on these dummy effects. In chart 15 to 17 the values of the estimated dummy coefficients for the volatility, average spread, and number of quotations are presented. They reveal an interesting feature. In the second half of the day BST time, from about the closing time of the European exchanges and until the closing time of the New York exchange, volatility is unusually high and low average spread. During this period there are few, or no, economic (or other public) announcements from Europe. Most US economic announcements are made before the opening of the New York Stock Exchange, at 13:30 BST. There is a small spike at the relevant half-hour (27), but this remains quite small compared with the higher volatilities apparent later on in the US market day. Hence, it seems that public news is not the explanation of this volatility increase. Furthermore, this increase seems even more difficult to explain in the light of the ADMATI/PFLEIDERER (1988) theory. During this period we certainly have a reduction in the number of traders in the market, as only the New York exchange is in operation; so this increase can hardly be attributed to an increase in the number of liquidity traders. Notice that the same phenomenon is also apparent in the Deutsch Mark and Yen/ $ markets (see DEMOS/GOODHART (1992)).

There is then a decrease in volatility until the early morning 0:30-1:00 BST interval, just before the Japanese economic related news are released (at 1:00 BST), when volatility is at its highest point, and average spread and activity are still close to their norms. We are not, however, able to explain why volatility should peak at this time. Following that we have an increase in activity and average spread and a dramatic decrease in volatility, which, except from the Japanese lunch time, stays at its norm level till the opening of the American markets.

Charts 15-17: Values of Estimated Dummy Coefficients: Volatility, Average Spread, and Number of Quotations

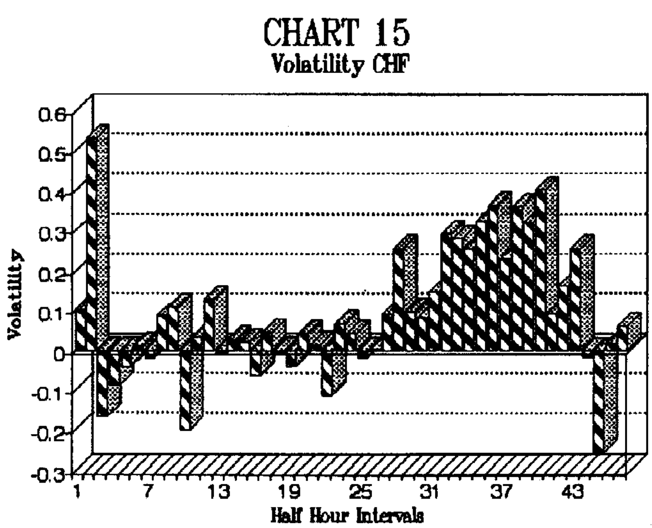

CHART 16
Average Spread CHF

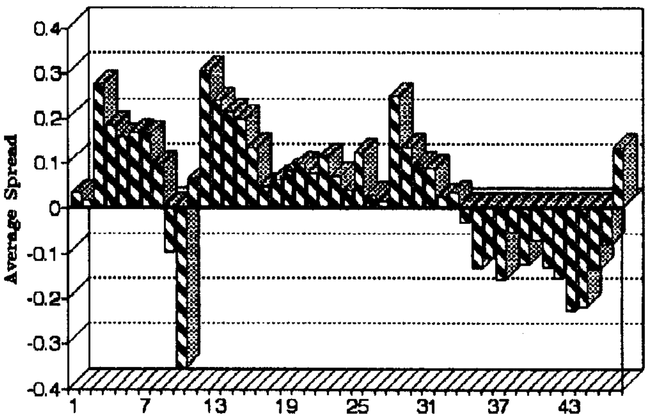

CHART 17
Number of Quotations CHF

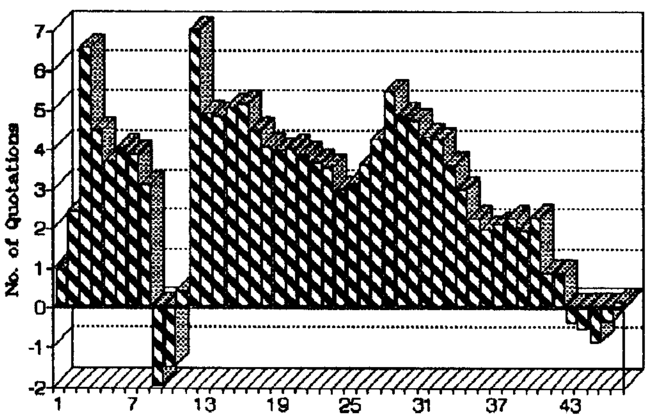

Table 12: Estimated Weekly and Daily Dummy Variables in Equations (4)-(6)

Parameter	Volat. Equat.	Sprea. Equat.	Quot. Equat.
Const.	-0.773	1.366	0.899
	(-2.72)	(7.52)	(0.90)
Weak 1	-0.207	-0.075	-0.415
	(-4.70)	(-2.05)	(-2.61)
Weak 2	-0.088	-0.132	-0.306
	(-2.09)	(-4.58)	(-2.18)
Weak 3	-0.130	-0.098	-0.188
	(-3.33)	(-3.35)	(-1.39)
Weak 4	-0.093	-0.074	-0.180
	(-2.52)	(-2.90)	(-1.52)
Monday	-0.061	-0.018	0.293
	(-1.66)	(-0.72)	(2.48)
Tuesday	-0.091	-0.009	0.429
	(-2.39)	(-0.35)	(3.17)
Wednes.	-0.046	-0.006	0.404
	(-1.36)	(-0.22)	(3.73)
Thursd.	-0.021	0.022	0.402
	(-0.65)	(0.87)	(3.82)

Robust to Heteroskedasticity t-statistics in parenthesis

There are signs of a significant pattern in volatility between days of the week (see Table 12) with the highest volatility on Friday, which is also observed in the forex future data (see Harvey and Huang (1991)), and least on Tuesday. The average spread however shows no signs of such a pattern. The weekly dummies, during this period, showed a pattern of increase in market volatility, but again this pattern is not clear for the average spread.

6. Conclusions

The relationship between the number of entries over Reuters' screen and the size of the market, as represented by turnover in the spot market, is subject to several major biasses, as emphasized througout this paper. Deals are arranged through brokers, or using other rival information systems, and the share of business done via Reuters may vary considerably between centres and between banks, but unobservably.

Where our data, we would claim, are reasonably representative, reliable and original occurs when we chart the time pattern of activity in, and through, each day in aggregate, by market area, centre, bank and branch. We believe that there is a close correlation between the frequency of price entry and turnover in the market, though we cannot prove, or disprove, the latter (yet), since data series for the latter are not (publicly) available to allow us to do so.

In this respect the daily pattern of price entries, during the weekdays, for the Swiss Franc/$ is closely similar to that of the Dm/$ and £/$. The number of the entries for CHF is slightly less than for Yen, but more than for £. There were a few more financial centres from which bank branches entered quotes for CHF, than for Yen and £, but, on the other hand, there were fewer bank branches and banks doing so, so the various measures showed a higher degree of market concentration, than for Yen or (just) for £. Overall, on the evidence of this data set, the Swiss Franc/$ and £/$ markets are closely similar in general format, as measured by the pattern of the number of quotations, though the £/$ market is even more concentrated in London than the CHF market in Zurich.

The contemporaneous relationship between the number of quotations, a measure of volatility, and spread is examined in part V. The number of the quotations plays an important role in the determination of volatility and average spread, either directly or through the error term. The contemporaneous correlation between this variable and volatility leads us to hypothesise that the former process could be a proxy for the volume of trade, or the number of transactions in the forex market, for which data are unavailable. This is in line with studies in stock market volume and volatility data (see GALLANT/ROSSI/TAUCHEN (1990) and LAMOUREUX/LASTRAPES (1990)). An important result of the above theories, for the stock market, and this paper, for the forex market, is that the contemporaneous determination of volatility and volume considerably reduces the GARCH type effects. Moreover, our results indicate that the inclusion of half-hourly dummies further reduce the GARCH type effects in the conditional variance of the exchange rate. What remains of such GARCH effects can then probably be attributed to private information and the uncertainty associated with it.

Informational theories can only partially explain the facts documented here. Although, high trading and volatility at the opening of markets can be explained along the lines of the ADMATI/PFLEIDERER (1988) theory, the different behaviour of the exchange rate in different markets at the same (and different) time periods points towards the need to take into account local and currency specific behaviour. Finally, having fitted weekly, daily and half-hour dummies, we can identify inter and intra-day patterns of activity, volatility and average spread. Some of these, eg the impact of the Tokyo lunch hour, we have documented in section II. Others are already well known in markets, eg the rise in spreads and decline in activity on Fridays. But we were surprised by the finding of the continuing high volatility, throughout the period of the US market, despite steadily falling activity, which we had expected. Much of the public information on economic news in the USA is released at, or before, the market opening, so exactly what keeps volatility so high during the afternoons in the USA is a mystery to us.

Appendix 1

The foreign exchange market uses electronic screens, e.g. Reuters, Telerate, Knight Ridder, to transmit indicative (i.e. not contractually binding) information about the prices that the main players, i.e. the large international banks, would be prepared to buy US$s in exchange for Swiss Francs, the bid price, or to sell US$s in exchange for Swiss Francs, the ask price. All spot currencies are traded bilaterally with $; cross rates are derived (it can be done electronically and automatically) from the basic bilateral $ rates. In all cases but four the rates are quoted in units of home currency, in this case Swiss Franc, per $; in the case of £, Irish punt, Australian $ and Ecu, the quotations are in terms of $s per unit of domestic currency.

The data on the FXFX arrive as follows (with thanks to Mark Jones and others at Reuters for their assistance). Each bank on the Reuters network has its own page (and associated mnemonic) allocated to it, which it updates itself, as and when it decides to do so. These pages are continuously scanned by Reuters. Whenever a bank revises a spot quote of a currency on FXFX, (or the minor currencies exhibited on FXFY), this latest quote is flashed up on the main FXFX(Y) screen. There are some security/checking gateways, but the process is almost instantaneous (i.e. hardly a second). During busy periods, which will be defined further below, the screen shows extremely frequent, rapid changes, as new entries come in. If, two banks should revise a quote for the same spot currency virtually simultaneously, i.e. bank Y revises at the same identical instance as bank X's quote is going through, then bank Y's revision will never appear on the screen: the transmission system is, however, so quick that no queuing arrangement has been thought necessary.

In order to analyse the process, some organization of the data is necessary. Using hardware kindly made available by Reuters and software developed by Philip Shearer and operated by Dr Russell Lloyd, we transformed the data coming onto the screens into time series separately for each currency.

Footnotes

[1] On a very few occasions prices are incorrectly input, and reach the screen, despite Reuters checks. We designed two filters to catch such mistakes. First an erroneous input, e.g. with the order of the bid/ask reversed, such as 1.7010/00 rather than 1.7000/10, would often show up in the form of an implausibly large spread. Second, an incorrect input would often show up as an implausibly large jump, followed by a reversal, in the data. We checked all outlying spread/jump observations, and threw out the implausible inputs, (averaging about 1 rejection per 400 inputs). The data examined here are from the filtered series.

[2] BANK FOR INTERNATIONAL SETTLEMENTS (1990), p. 3.

[3] This can be demostrated since Reuters AAMM "news" pages are part of our data set (see also HARVEY/HUANG (1991)).

[4] For space economy the Saturday and Sunday Charts are not presented.

[5] This however should not be confused with the number of entries as a measure of relative intradaily or weekly activity.

References

ADMATI, A.R. and P. PFLEIDERER, (1988): "A Theory of Intraday Patterns: Volume and Price Variability", The Review of Financial Studies 1, pp. 3-40.

BANK FOR INTERNATIONAL SETTLEMENT (1990): "Survey of Foreign Exchange Market Activity", Basel.

BASMANN, R. (1974): "Exact Finite Sample Distribution for Some Econometric Estimators and Test Statistics: A Survey and Appraisal", in: Intriligator, M. and D. Kendrick (Eds.): Frontier of Quantitative Economics Volume 2, Chapter 4, North Holand.

BERKMAN, H. (1991): "The Market Spread, Limit Orders and Options", Report No. 9007, Erasmus University, Rotterdam.

BOLLERSLEV, T., R.Y. CHOU and K.F. KRONER (1992): "Arch Modelling in Finance", Journal of Econometrics 52, pp. 5-59

DASSIOS, A. and A. DEMOS, (1991): "Some Econometric Tests on Real Time Forex Market Data", Discussion Paper in Financial Economics FE-5/91, Birkbeck College, University of London.

DEMOS, A. and C. GOODHART, (1992): "The Interaction between the Frequency of Market Quotations, Spread, and Volatility in the Foreign Exchange Market", Financial Markets Group Discussion Paper No. 152, London School of Economics.

ENGLE, R., T. ITO, and W. LIN (1990): "Meteor Showers or Heat Waves? Heteroskedastic Intra-daily Volatility in the Foreign Exchange Market", Econometrica 58, pp. 525-542.

ENGLE, R. and V. NG (1991): "Measuring and Testing the Impact of News on Volatility", mimeo University of California.

THE ECONOMISTS ESSENTIAL DIARY INFORMATION (1989).

FOSTER, F. and S. VISWANATHAN, (1990): "A Theory of the Individual Variations in Volumes, Variances and Trading Costs in Securities Markets", Review of Financial Studies 3, pp. 593-624.

GALLANT, A., P. ROSSI and G. TAUCHEN (1990): "Stock Prices and Volume", mimeo Dept. Economics, Duke University.

GOODHART, C. and A. DEMOS (1990): "Reuters Screen Images of the Foreign Exchange Market: The Deutschemark/Dollar Spot Rate", The Journal of International Securities Markets 4, pp. 333-348.

GOODHART, C. (1990): "'News' and the Foreign Exchange Market", London School of Economics, Financial Markets Group Discussion Paper No. 71.

GOODHART, C. and A. DEMOS (1991): "Reuters Screen Images of the Foreign Exchange Markets: The Yen/Dollar and Sterling/Dollar Spot Market", The Journal of International Securities Markets 5, pp. 35-64.

GOODHART, C. and M. GIUGALE (1991): "From Hour to Hour in the Foreign Exchange Market", Manchester School Journal (forthcoming).
HARVEY, C. and R. HUANG (1991): "Information and Volatility in the FX Markets", Vanderbilt University Working Paper.
HAUSMAN, J. (1978): "Specification Tests in Econometrics", Econometrica 46, pp. 1251-1271.
HO, T. and H. STOLL (1983): "The Dynamics of Dealer Markets Under Competition", Journal of Finance 38, pp. 1053-1074.
ITO, T. and V. ROLEY (1987): "News from USA and Japan: Which moves the Yen/Dollar Exchange Rate?", Journal of Monetary Economics 19, pp. 255-277.
JONES C., G. KAUL and M. LIPSON (1991): "Transactions, Volumes and Volatility", mimeo University of Michigan, School of Business Administration.
LAMOUREUX, C. and W. LASTRAPES (1990): "Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects", Journal of Finance 45, pp. 221-229.
LAUX, P. and L. NG (1991): "Intraday Heteroskedasticity and Comovements in the Foreign Currency Futures Market", mimeo Dept. Finance, Univarsity of Texas.
WASSERFALLEN, W. and H. ZIMMERMANN (1985): "The Behaviour of Intra-daily Exchange Rates", Journal of Banking and Finance 9, pp. 55-72.
WASSERFALLEN, W. (1989): "Flexible Exchange Rates: A closer look", Journal of Monetary Economics 23, pp. 511-521.
WU, D. (1973): "Alternative Tests of Independence between Stochastic Regressors and Disturbances", Econometrica 41, pp. 733-750.