

From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets*

**Dominique M. Guillaume¹, Michel M. Dacorogna², Rakhil R. Davé²,
Ulrich A. Müller², Richard B. Olsen², Olivier V. Pictet²**

¹ Financial Markets Group, London School of Economics and C.S.A.E. Institute of Economics and Statistics, University of Oxford, St. Cross Building, Manor Road, Oxford OX1 3UL, United Kingdom (e-mail: dominique.guillaume@economics.ox.ac.uk)

² Olsen & Associates, Research Institute for Applied Economics, CH-8008 Zürich, Switzerland

Abstract. This paper presents stylized facts concerning the spot intra-daily foreign exchange markets. It first describes intra-daily data and proposes a set of definitions for the variables of interest. Empirical regularities of the foreign exchange intra-daily data are then grouped under three major topics: the distribution of price changes, the process of price formation and the heterogeneous structure of the market. The stylized facts surveyed in this paper shed new light on the market structure that appears composed of heterogeneous agents. It also poses several challenges such as the definition of price and of the time-scale, the concepts of risk and efficiency, the modeling of the markets and the learning process.

Key words: Foreign exchange rates, intra-daily, heterogeneous, distribution

JEL classification: F31, G15, E44

Mathematics Subject Classification (1991): 90-02, 62-07, 62P20

* T. Bollerslev, P. Dall'Acqua, P. De Grauwe, C. Gauvreau, G. Gielens, Ch. Goodhart, Ph. Hartmann, M. Melvin, P. Korndörfer, R. Ward and the co-editor A. Shiryayev provided useful comments. The first author gratefully acknowledges financial support from the Belgian National Foundation for Scientific Research N.F.W.O. contract 2.0065.94.

Editorial Note:

This paper is of a survey and descriptive character, embracing a broad range of different aspects of empirical financial studies. The financial data provide "raw material" for decision making as well as for corrections of existing theoretical models. With the increasing availability of high frequency data a deeper analysis of the statistics of financial time series becomes of growing importance. In publishing this article, the Editorial Board pursues the aim of stimulating further research in this direction.

Manuscript received: May 1996; final version received: September 1996

1. Introduction

Since the beginning of the 1990s, academic researchers have been gaining new insights into the behavior of the foreign exchange (FX) markets by analyzing intra-daily data. Indeed, much information about the FX markets, both quantitative and qualitative, was not used up to now. Daily data which were much used in the 80's, represent only a very small subset of the information available at intra-daily frequencies as they really are only the average of five intra-daily quoted prices of the largest banks around a particular time. The number of data points available intra-daily is 100 to 1000 times larger. Furthermore, the actual news or information at the disposal of the traders with the corresponding time stamp is displayed along with FX rates quotations on the data vendors screens. Finally, in addition to the volatility and the spread, new quantities for the description of the price evolution like the tick frequency or the directional change frequency can also be directly computed from the more comprehensive data set.

On the basis of this information set, there is now a rapidly growing body of empirical studies on the behaviour of the intra-daily FX markets. This is especially interesting in that it opens new paths for understanding the behaviour of financial markets and suggests the broadening of concepts such as risk or market efficiency. The analysis of intra-daily data may also help clarify the intuitive approach of intra-day traders, whose operations account for more than 75% (Bank of International Settlements, 1993) of the FX market volume. Indeed, as we shall see in part 4, the heterogeneous structure of intra-daily data may explain the success of technical analysis so much in favour amongst intra-day traders since these rules try to intuitively take advantage of the interaction of different components of the markets [Dunis and Feeny 1989; Neftci 1991; Surajaras and Sweeney 1992; Taylor and Allen 1992; Pictet et al. 1992; Levich and Thomas 1993]. Finally, this set of facts about the intra-daily FX rates sheds some new light on different modelling approaches to the FX market. The purpose of this survey is to bring these facts together in order to give an overall picture of the FX markets and to provide a common basis for future empirical research and theoretical modelling.

Besides extending the current literature in that it is the first review of the *intra-daily* foreign exchange markets, this paper also offers some innovative ideas¹. First, it proposes definitions of the variables to be studied and points out their limitations. Second, we distinguish between the general characteristics of intra-daily data and those of the process of price formation that takes place at the highest frequency (within the 10 to 15 minutes time-interval) in order to highlight the specificity of this process. Third, different major issues such as the characterisation of the distribution of the price change, the definition and the use of the time-scale, the concept of risk or market efficiency, the process of learning and the modelling approaches to the market will be tackled in this survey.

Our focus is on the nominal spot exchange rates of major currencies. In particular we do not study well-known relationships between spot rates and other financial variables such as forward rates or interest rates which have not yet been

investigated with intra-daily data. Experience has indeed shown that known and well accepted empirical regularities of daily or weekly data do not always hold up in intra-daily analysis. Moreover, the different structures of these financial markets may play an important role for the analysis of their statistical properties. We have also not systematically investigated the relationship between exchange rates and major macro-economic variables although new insights on how to tackle this issue are presented.

As will be shown, looking at intra-daily data, the homogeneity of market agents apparent at lower, weekly or daily, frequencies disappears. A new wealth of structure is uncovered that demonstrates the complexity of the FX market. This complexity can be explained by the interaction of market agents with heterogeneous objectives resulting from their different geographical location, the various types of their institutional constraints and their risk profiles. This is the key for future modelling of the dynamics of the markets.

The remainder of this paper is organised as follows. Section 2 briefly describes the functioning of the foreign exchange markets and the data. Section 3 proposes a set of definitions for the variables under study. Section 4 regroups the set of empirical regularities under three different topics: the distribution aspects of the price process, the process of price formation and the heterogeneous structure of the markets. Section 5 concludes the paper.

2. Description of the FX market and the data set

The usual description of the FX markets made by international organisations such as the Bank for International Settlement or the International Monetary Fund emphasises the presence of different geographical markets and different types of agents (BIS 1993; International Monetary Fund 1993). However, this essential fact was not very apparent from the inspection of daily or weekly data and its implications were seldom considered in theoretical modelling. In this section, we shall summarise the essential descriptive knowledge we have of the FX market. Additionally, we describe the main source and type of data available for this market. As we shall see, the information content of intra-daily data is much richer than that of daily or weekly data. However, the extent to which this information set can be increased and the degree of precision one can obtain are limited. This is reflected in the distinct characteristics of extremely high frequency data corresponding to the price formation process.

The FX market is a 24 hours global market, mostly inactive during weekends and during national holidays. The first observation of the week arrives at 22:30 Greenwich Mean Time (GMT) on Sunday with the opening of the Asian markets and the last observation comes from the West Coast of the USA at about 22:30 (GMT) on Friday². Although the FX market is virtually global through its electronic linkages, its activity pattern can be divided into three continental components: East Asia with Tokyo as major trading center, Europe with London as major trading center and America with New York as major trading center³

(Goodhart and Demos 1990). Except for the four major currencies against the USD (DEM, JPY, GBP and CHF)⁴ currencies tend to be traded more specifically in their own geographical markets. As we shall see in Sect. 3.C below, both these global and local characteristics of the FX markets are reflected in the statistical properties of the data.

The FX market is also the largest financial market with a daily turn-over⁵ of USD 832 billion in April 1992 (BIS 1993): that is, more than the total non-gold reserves (USD 555.6 billion) of all industrial countries in 1992 (International Monetary Fund 1993) and more than the triple of the turnover in 1986. This fast growing volume of transactions is increasingly made up of short term, intra-daily, transactions and results from the interaction of traders with different time-horizons, risk-profiles or regulatory constraints. On the one side, the non-financial institutions such as non-financial corporations, institutional investors (mutual funds, pension funds, insurance companies) and hedge funds⁶ are shifting their FX activities from long term (buy and hold) investment to short term (profit-making) transactions. This movement is both enabled and enhanced by the development of real-time information systems and the decrease of transaction costs following the liberalisation of cross-border financial flows. On the other side, this flow of short and long term transactions initiated by non-financial institutions on the retail market is the origin of an even larger – by a factor of four to five times – flow of intra-daily transactions between the dealers (the 50 largest banks and a few securities houses) on the wholesale market. These dealers who are not usually allowed to take overnight positions, move to reduce their risk with each other (Lyons 1994). Still on the wholesale market, but in contrast with other players, central banks can afford relatively large open positions and can thereby have a significant impact on the market in the long run. These different types of traders can of course be found within the same company or the same type of institution⁷. Again, their presence is largely reflected in the statistical characteristics of the data as we shall see in Sect. 3.

Although the FX market consists of spot and longer term instruments, we concentrate here on the spot market⁸ that operates mainly through electronic screens from financial news agencies like Reuter's, Knight Ridder or Telerate, where market makers enter quotes for their bid and ask prices. The actual deals are then made over the telephone. Transactions can also be made directly through automated dealing systems that offer the advantage of displaying actual transaction prices and volumes. Despite the growing importance of such markets⁹, these transaction data became only very recently available to researchers in a quite limited amount¹⁰. As will be stated in Sect. 4.2, the properties of real transaction data do not differ from those of Reuter's FXFX data except for frequencies higher than 10 to 15 minutes. Therefore this survey concerns itself with tick-by-tick quoted prices from Reuter's FXFX page and equivalent service from Knight Ridder and Telerate¹¹. To give an indication of the largest database of such quotes currently existing, Table 1 displays the number of quotes of the Olsen&Associates database on which some of the studies of this survey including the following papers, were made. Table 1 also gives the average number of quotes available per FX rate.

On the largest market, the USD/DEM, 4,500 quotes per day are available; that is, there is an average of 3 to 4 new quotes per minute, but this average can rise to 15-20 quotes per minute during the busiest periods.

Table 1. Data set. Number of data for the main FX rates against the USD (period 01.01.1987-31.12.1993) and the DEM (period 01.10.1992-31.12.1993) using the filter for outliers of Dacorogna et al. (1994)

Rate	Number of data	Businessday frequency
USD_DEM	8,238,532	4,500
USD_JPY	4,230,041	2,300
GBP_USD	3,469,421	1,900
USD_CHF	3,452,559	1,900
USD_FRF	2,098,679	1,150
JPY_DEM	190,967	630
FRF_DEM	132,089	440
ITL_DEM	118,114	390
DEM_GBP	96,537	320
NLG_DEM	20,355	70

An example of Reuter's FXFX page which displays constantly updated quotes from the banks that subscribe to its service is given in Table 2. A quoted price of 1.6290/00 for the USD-DEM rate expresses the willingness of the market maker to buy USD at 1.6290 DEM, and sell USD at 1.6300 DEM. Actual trading prices and volumes are not known from this page. However reputation considerations prevent market makers from quoting prices at which they would actually not be willing to trade. Therefore real transaction prices will usually tend to be comprised within the quoted bid/ask spread (Petersen and Fialkowski 1994).

In contrast to daily or weekly data, collecting these tick-by-tick quotes presents a number of practical problems such as transmission delays and breakdowns or aberrant quotes due to human and technical errors. Therefore, it is important to implement filters to eliminate outliers. The reader can find examples of such filters in Goodhart and Figliuoli (1991) and Dacorogna et al. (1993, 1994). As noted in Dacorogna et al. (1994), the percentage of quotes that are eliminated by such filters is very low (less than 0.5%).

3. Definition of the variables of interest

Adequate analysis of such intra-daily data relies on an explicit definition of the variables under study. Some of these definitions are *redefinitions* of variables in current everyday use, it must be stressed. These include the price, the change of price, the volatility and the spread. The remainder are newly developed variables, formulated to better capture the peculiarities of the intra-day market. These include

Table 2. Reuter's FXFX page screen. The first column gives the time (for example, for the first line, '07:27 '), the second column gives the name of the currency ('DEM/USD'), the third column gives the name of the bank subsidiary which publishes the quote given as a mnemonic ('RABO' for the Rabobank), the fourth column gives the name of the bank ('Rabobank'), the fifth column gives the location of the bank as a mnemonic ('UTR' for Utrecht), the sixth and seventh column give the bid price with 5 digits ('1.6290') and the two last digits of the ask price ('00'), the last two columns give the highest ('1.6365') and the lowest ('1.6270') quoted prices of the day

0727	CCY	PAGE	NAME	*	REUTER	SPOT	RATES	*	CCY	HI*EURO*LO	FXFX
0727	DEM	RABO	RABOBANK	UTR	1.6290/00	*	DEM	1.6365	1.6270		
0727	GBP	MNBX	MOSCOW	LDN	1.5237/42	*	GBP	1.5245	1.5207		
0727	CHF	UBZA	U B S	ZUR	1.3655/65	*	CHF	1.3730	1.3630		
0727	JPY	IBJX	I.B.J	LDN	102.78/83	*	JPY	103.02	102.70		
0727	FRF	BUEX	UE CIC	PAR	5.5620/30	*	FRF	5.5835	5.5582		
0726	NLG	RABO	RABOBANK	UTR	1.8233/38	*	NLG	1.8309	1.8220		
0727	ITL	BCIX	B.C.I.	MIL	1592.00/3.00	*	ITL	1596.00	1591.25		
0727	ECU	NWNT	NATWEST	LDN	1.1807/12	*	ECU	1.1820	1.1774		

XAU	SBZG	387.10/387.60	*	ED3	4.43/ 4.56	*	FED	PREB	* GOVA	30Y	
XAG	SBCM	5.52/ 5.53	*	US30Y	YTM 7.39	*	4.31-	4.31	*	86.14-15	

the tick frequency, the volatility ratio and the directional change frequency. All are necessary for understanding the Stylized Facts and we also propose this set as a unifying reference for future studies. An extensive notation is given along with the usual simplified notation to make all the underlying parameters explicit. For each variable, we first give a formal definition and then discuss its use.

3.1. Definition 1: the price

The *price* at time τ_j , $x(\tau_j)$, is defined as

$$x(\tau_j) \equiv [\log p_{bid}(\tau_j) + \log p_{ask}(\tau_j)] / 2 \quad (1)$$

where $\{\tau_j\}$ is the sequence of the tick recording times which is unequally spaced. An alternative notation is

$$x(t_i) \equiv x(t_i, \Delta t) \equiv [\log p_{bid}(t_i) + \log p_{ask}(t_i)] / 2 \quad (2)$$

where $\{t_i\}$ is the sequence of the regular spaced in time data and Δt is the time interval ($\Delta t = 1$ day, $\Delta t = 1$ hour, ...).

Definition 1 takes the average of the bid and ask price rather than either the bid or the ask series as a better approximation of the transaction price. Indeed as Fact 6 will point out, market makers frequently skew the spread towards a more favorable price to offset their position. In that context, the bid (or ask) price acts as a dummy value. Furthermore, in Definition 1, the average of the logarithms of the bid and ask prices rather than the logarithm of the average is taken since

the former quantity has the advantage of behaving anti-symmetrically when the price is inverted.

One important issue in the case of intra-daily data is the use of the right *time-scale*. Contrary to daily and weekly data, tick-by-tick data are indeed irregularly spaced in time, τ_j . However, most statistical analyses rely upon the use of data regularly spaced in time, t_i . For obtaining price values at a time t_i within a data hole or in any interval between ticks we use the linear interpolation between the previous price at τ_{j-1} and next one at τ_j , with $\tau_{j-1} < t_i < \tau_j$. As advocated in Müller et al. (1990), linear interpolation is the appropriate method for interpolating in a series with independent random increments for most types of analyses. An alternative interpolation method might be to use the most recently published price as in Wasserfallen and Zimmermann (1985) although this introduces an inevitable bias in the data. However, as long as the data frequency is low enough, the results do not depend too much on the choice of either method. Although regularly time spaced data are used in most of the definitions below, irregularly time spaced data could alternatively be used by replacing t_i by τ_j . Finally, in addition to these two time-scales, other time-scales have been proposed to model characteristics of the intra-daily FX market such as the seasonality (Dacorogna et al. 1993), the heteroskedasticity (Zhou 1993) or both of them (Müller et al. 1993; Guillaume et al. 1996).

Another important issue is the definition of an *effective price*. Indeed, at frequencies higher than 10 minutes, the size of the spread is of the same order of magnitude as the size of the price changes. Moreover, the quoted spread does not exactly reflect the real spread, which is usually smaller as reported in Goodhart et al. (1994)¹². Furthermore, because of transmission delays, it may be, for example, that market maker B enters a quote after market maker A, but that the quote of market maker B is the first to appear on the multi-contributor page of Reuters. Similarly, holes due to transmission breakdowns become more significant at such frequencies. Therefore, at frequencies higher than approximately 10 minutes, reliable analysis of prices can only be done on the basis of real transaction prices or with a thorough knowledge of the shortcomings of the database.

In the absence of such real transaction prices, we may define an effective price algorithm by looking at the properties of the prices and the market structure organization. Assuming that quotes have a life-time of approximately 2 minutes during periods of average activity, one could take the best bid and ask quotes available in such a time window or the averages of the bids and of the asks. Another idea for such an algorithm would be to eliminate the negative first-order autocorrelation of the prices present at such frequencies (see Fact 5). An example of an algorithm for the computation of effective price is given in Bollerslev and Domowitz (1993) where the trade-matching algorithm designed for the interbank market by Reuters – Dealing 2000 system – is used. Interestingly, the prices generated by this algorithm exhibit a positive rather than negative first-order autocorrelation. In contrast, Goodhart et al. (1994) still obtain a negative first-order autocorrelation - although less pronounced - in their analysis of the Dealing 2000-2 system from Reuters.

3.2. Definition 2: the change of price

The *change of price* at time t_i , $r(t_i)$, is defined as

$$r(t_i) \equiv r(\Delta t; t_i) \equiv [x(t_i) - x(t_i - \Delta t)] \quad (3)$$

where $\{x(t_i)\}$ is the sequence of equally spaced in time logarithmic price, and Δt is the fixed time-interval (10 minutes, 1 hour, 1 day, ...).

The change of the logarithmic price is often referred to as “return”. It is usually preferred to the price itself as it is the variable of interest for traders maximizing short term investment returns. Furthermore, its distribution is more symmetric than the distribution of the price. Finally, it is usually advocated that contrary to the price process which is clearly non-stationary, the process of the price changes should be stationary. Although unit root tests such as the Augmented Dickey Fuller test cannot reject the hypothesis of stationarity in the mean (Goodhart and Figliuoli, 1991), the autocorrelation of the volatility would probably cause the rejection of covariance stationarity by a test such as the one proposed in Loretan and Phillips (1994).

3.3. Definition 3: the volatility

The *volatility* at time t_i , $v(t_i)$, is defined as

$$v(t_i) \equiv v(\Delta t, S; t_i) \equiv \frac{1}{n} \sum_{k=1}^n |r(\Delta t; t_{i-k})| \quad (4)$$

where S is the sample period on which the volatility is computed (for example 1 day or 1 year) and n is a positive integer with $S = n\Delta t$. A usual example is the computation of the daily volatility as the average daily volatility over one year ($S = 1$ year, $n = 250$ and $\Delta t = 1$ day). Note that in an autocorrelation study, for example, only 1 data point may be used ($S = \Delta t$) or n might be multiplied by a factor of 2 or 3 if overlapping data are taken as suggested in Müller (1993).

In Definition 3, the absolute value of the returns is preferred to the more usual squared value or more generally to any power ε ($\varepsilon \in \mathbb{R}_0^+$) of $|r(\Delta t; t_i)|$. This is because the former quantity better captures the autocorrelation and the seasonality of the data (Taylor 1988; Müller et al. 1990; Grange and Ding 1993). This greater capacity to reflect the structure of the data can also be easily derived from the non-existence of a fourth moment in the distribution of the price changes (see Fact 4).

Although Definition 3 is the most appropriate for the assessment of risk or for forecasting, one might prefer other definitions of the volatility that give more weight to the tails of the distribution: for instance the cube root of the third moment for the evaluation of extreme downside risk in portfolio optimisation as in Roy (1952). One could also prefer the use of implicit volatility such as defined by the option model (Cox and Rubinstein 1985) or the Generalized Autoregressive Conditional Heteroskedastic (GARCH) model (Bollerslev 1986). However,

although both approaches might appear appealing at the daily frequency, their use at the intra-daily frequency presents important drawbacks; on the one side, the implicit volatility cannot be computed at very high frequency since options are not quoted at such frequencies; on the other side, as a consequence of the heterogeneity of the FX market, intra-daily FX rates cannot be described by one homogeneous GARCH model (Guillaume 1994).

3.4. Definition 4: the relative spread

The relative *spread* at time t_i , $s(t_i)$, is defined as

$$s(t_i) \equiv \log p_{ask}(t_i) - \log p_{bid}(t_i) \quad (5)$$

Definition 4.1:

The *log spread* at time t_i , $\log s(t_i)$, is defined as

$$\log s(t_i) \equiv \log(\log p_{ask}(t_i) - \log p_{bid}(t_i)) \quad (6)$$

In the above definition, the relative spread $s(t_i)$ is preferred to the nominal spread ($p_{ask}(t_i) - p_{bid}(t_i)$) since it is dimensionless and can therefore be directly compared between different currencies. The spread of the inverse rate (e.g. JPY per USD instead of USD per JPY) is simply $-s(t_i)$, so that the variance of $s(t_i)$ is invariant under inversion of the rate.

It might however be sometimes preferable to study $\log s(t_i)$ instead of $s(t_i)$, because the former quantity has a more symmetric behavior between low and high values. The relative spread is indeed a positive bounded quantity that has a skewed distribution function. It cannot be much lower than the typical value – that is, 10, 7, 5, or 15 basis points – but, on the positive side, it can exceed the typical value by a factor of 2 or more. Moreover, these high spreads which tend to numerically dominate a statistical analysis, usually occur in thin and unimportant markets, for example over the weekend or during the East Asian lunch break.

The spread is indicative of the transaction and inventory costs of the market maker who is under reputation consideration pressures. It is also affected by the degree of informational asymmetries and competitiveness. Thus, the spread depends both on the cost structure of the quoting bank and on the habits of the market. On the other side, it is the only source of cost for the traders since intra-daily credit lines on the foreign exchange markets are free of interest¹³.

3.5. Definition 5: the tick frequency

The *tick frequency* at time t_i , $f(t_i)$, is defined as

$$f(t_i) \equiv f(S; t_i) \equiv \frac{1}{S} N(\{x(\tau_j) \mid \tau_j \in (t_i - S, t_i]\}) \quad (7)$$

Definition 5.1:

The log *tick frequency* at time t_i , $\log f(t_i)$, is defined as

$$\log f(t_i) \equiv \log f(S; t_i) \quad (8)$$

where $N(\{x(t_j)\})$ is the counting function and S is the sample period on which the counting is computed. The alternative log form has been found to be more relevant in Demos and Goodhart (1992) .

The tick frequency is sometimes taken as a proxy for the transaction volume on the markets. As the name and location of the quoting banks are also given, the tick frequency is also sometimes disaggregated by bank. However, equating tick frequency to transaction volume or using it as a proxy for both volume and strength of bank presence suffers from the following problems: First, although it takes only a few seconds to enter a price quotation in the terminal, if two market makers happen to simultaneously enter quotes, only one quote will appear on the data collector's screen; Second, during periods of high activity, some operators may be too busy to enter the quote into the system; Third, a bank may use an automatic system to publish prices to advertise itself on the market. Conversely, well-established banks might not need to publish as many quotes on smaller markets; Fourth, the representation of the banks depends on the coverage of the market by data vendors such as Reuters or Knight Ridder. This coverage is changing and does not totally represent the whole market. For example, Asian market makers are not as well covered by Reuters as the Europeans. Asian market makers are instead more inclined to contribute to the more local financial news agencies such as Minex; Fifth, trading strategies of big banks are highly decentralized by subsidiary. Even between the back office and the trading room¹⁴ or within the trading room itself, different traders may have completely different strategies.

3.6. Definition 6: the volatility ratio

The *volatility ratio* at time t_i , $Q(t_i)$, is defined as

$$Q(t_i) \equiv Q(\Delta t, n; t_i) \equiv \frac{|\sum_{k=1}^n r(t_{i+k})|}{\sum_{k=1}^n |r(t_{i+k})|} \quad (9)$$

The volatility ratio defined above is simply a generalization of the variance ratio introduced in Lo and MacKinlay (1988) and Poterba and Summers (1988) where the absolute value of the price change instead of the variance is used as a measure of the volatility to take into account the statistical properties of the data (see Definition 3). The ratio can take values between "1" when the price changes follow a pure trend and "0" when they behave purely randomly.

The volatility ratio has been used in a variety of applications, including the effect of structural changes on prices, hypothesis testing in the empirical literature on the micro-structure of the markets and the identification of the nature of news. Here, however, we would like to stress the potential use of the volatility ratio as

a general statistic to measure the trend-following behavior of the price changes. This can be viewed as an alternative measure of the risk next to the volatility. Indeed, even though the two time series may have the same average volatility, the volatility ratio will be equal to “1” when the price changes follow a pure trend and to “0” when they behave randomly. Although the volatility ratio has many interesting features, we introduce a more precise statistic which seems to be more appealing to quantify the trend-following behavior of the price changes.

3.7. Definition 7: the directional change frequency

The *directional change frequency* at time t_i , $d(t_i)$, is defined as

$$d(t_i) \equiv d(\Delta t, n, r_c; t_i) \equiv \frac{1}{n\Delta t} N(\{k \mid m_k \neq m_{k-1}, 1 < k \leq n\}) \quad (10)$$

where $N(\{k\})$ is the counting function, $n\Delta t$ the sampling period on which the counting is performed, m_k indicates the mode – upwards or downwards – of the current trend and r_c is a threshold value used to compute the change of mode. The directional change frequency, $d(t_i)$, is simply the frequency of significant mode (m_k) changes with respect to the latest extremum value (max_k or min_k) and a constant threshold value r_c ¹⁵.

In contrast with the definition of the volatility where the time interval is the arbitrarily set parameter and the amplitude of the change of price is the varying parameter, in the above formulation, the time is varying and the threshold is fixed. Thus, the definition also takes into account gradually occurring directional changes.

The directional change frequency is similar to the volatility ratio defined above in that they both measure the trend-following behavior of the price changes and, as such, provide an alternative measure of the risk. However, unlike the volatility ratio, it is based on a threshold which is a measure of the risk quite natural to traders as put by one of them; “Although volatility can tell us the general environment of the market, we are actually more interested in the timing of our trades¹⁶. The knowledge of whether prices are likely to move more than a certain threshold allows us to decide when we need to close a position. The height of this threshold will vary according to our attitude towards risk.” The use of thresholds and measures of trends is also very familiar to chartists [see, for example, the technique of Point and Figure Charts in Meyers (1989)].

4. Stylized facts

We can now consider a new set of stylized facts describing the characteristics of the FX market, and most importantly, how it functions. These facts have been grouped to show the distributional properties, the price formation process and the time series properties.

4.1. Distribution

The variety of opinions about the distributions of FX price changes and their generating process is wide. Some authors claim the distributions to be close to Paretian stable (McFarland et al. 1982), some to Student distributions (Boothe and Glassmann 1987), some reject any single distribution (Calderon-Rossel and Ben-Horim 1982).

Instead of looking at the center of the distribution, an alternative way to characterize the distribution is to look at the tails. Most types of distributions can indeed be classified into three categories (de Haan 1990): (i) Thin-tailed distributions for which all moments exist and whose cumulative distribution function declines exponentially in the tails; (ii) Fat-tailed distributions whose cumulative distribution function declines with a power in the tails; (iii) Bounded distributions which have no tails. A nice result is that these categories can be distinguished by the use of only one parameter, the tail index¹⁷ α with $\alpha = \infty$ for distributions of category (i), $\alpha > 0$ for category (ii) and $\alpha < 0$ for category (iii).

The empirical estimation of the tail index and the variance of this tail index crucially depends on the size of the sample. Indeed, for a given sample size, on the one hand, using too many observations introduces a bias in the tail index as some of the observations do not belong to the tail anymore but are from the center of the distribution; on the other hand, using too few observations introduces an inefficiency in the estimation of the variance of the tail index. Therefore, the very large sample size available with intra-daily data ensures that enough tail observations are present in the sample.

An important result is that the tails of a fat-tailed distribution are invariant under addition although the distribution as a whole may vary according to temporal aggregation (Feller 1971). That is, if weekly returns are Student-t identically and independently distributed, then monthly returns are not Student-t distributed¹⁸. Yet the tails of the monthly return distribution are like the tails of the weekly returns, with the same exponent α .

Another important result in the case of fat-tailed distributions concerns the existence of the moments of the distribution. Let X be the observed variable, c a scale variable and α the tail index. From

$$E [X^k] = c \int_1^{\infty} x^{k-\alpha-1} dx \quad , \quad (11)$$

it is easily seen that only the first k -moments, $k < \alpha$, are bounded.

Finally, the tail index reflects the interaction of different agents on the markets. Indeed, the probability of extreme events depends on the presence or absence of certain market participants such as medium-term investors or pure speculators due to changing market conditions. The mechanisms leading to these fat tails can be understood as follows: whenever (relatively) long term traders become active following some news or perturbation, shorter term traders become even more active and tend to reinforce the longer term fluctuations, thereby creating fat tails (see Fact 12). This is a kind of "cascade" effect where getting over a

particular threshold at one (long) time interval attracts exceedings of thresholds at other time horizons (see the fractal law in Fact 10). This cascade effect thus can be explained by models of imperfect, though rational, information aggregation. In the absence of common information on the other traders' preferences or beliefs on the impact of news, this will cause some learning mechanism by the intra-day dealers before the prices fully incorporate the new information. Here, we posit that the larger the tail index, the less friction in the adjustment of prices to external shocks and the more efficient the market will be. Indeed, since most fundamentals have thin tails, a small tail index can only result from large fluctuations due to the trading process.

Table 3. Estimated tail exponent. Estimated tail exponent α and its standard error for the main FX rates against the USD and some of the main (computed) cross-rates against the DEM. The results are taken from Dacorogna et al. (1994). The bias was estimated using a bootstrap method. In contrast to quoted cross-rates, computed cross-rates are obtained via the two bilateral rates against the USD. Their spread is thus approximately twice the normal spread

Rate	10m	30m	1h	6h
USD_DEM	3.11 \pm 0.33	3.35 \pm 0.29	3.50 \pm 0.57	4.48 \pm 1.64
USD_JPY	3.53 \pm 0.21	3.55 \pm 0.47	3.62 \pm 0.46	3.86 \pm 1.81
GBP_USD	3.44 \pm 0.22	3.52 \pm 0.46	4.01 \pm 1.09	6.93 \pm 10.79
USD_CHF	3.64 \pm 0.41	3.74 \pm 0.82	3.84 \pm 0.77	4.39 \pm 4.64
USD_FRF	3.34 \pm 0.22	3.29 \pm 0.47	3.40 \pm 0.69	4.61 \pm 1.21
FRF_DEM	3.11 \pm 0.41	2.55 \pm 0.23	2.43 \pm 0.23	3.54 \pm 1.42
DEM_NLG	3.05 \pm 0.27	2.44 \pm 0.08	2.19 \pm 0.12	3.37 \pm 1.43
DEM_ITL	3.31 \pm 0.51	2.93 \pm 1.17	2.54 \pm 0.49	2.86 \pm 0.98
GBP_DEM	3.68 \pm 0.35	3.63 \pm 0.42	4.18 \pm 1.17	3.22 \pm 0.79
DEM_JPY	3.96 \pm 0.41	4.18 \pm 0.90	4.13 \pm 1.05	4.71 \pm 1.61

4.1.1 Fact 1: Non stable, fat-tailed distribution. From Table 3, it appears that the main FX rates against the USD have a tail index of ≈ 3.5 , whereas the tail index for the EMS computed rates against the DEM has a lower value of ≈ 2.7 and the tail index for the DEM/JPY and GBP/DEM is even larger than for free-floating rates against the USD. Quoted cross-rates in the EMS on the other hand have an even smaller value, reflecting the upward bias of the results with computed cross-rates at the 10 minutes interval due to the larger spread (Dacorogna et al. 1994). This has the effect of gaussian noise. However, the error bars of the tail index on quoted cross-rates are much larger because of the small size of the sample for that exercise. The smaller tail index for these rates thus indicates that the reduced variance induced by the EMS set-up is at the cost of a larger probability of extreme events.

Table 3 indicates that FX rates belong to the class of fat-tailed non-stable distributions which have a finite tail index. Furthermore, estimations of the kurtosis of the FX rate returns for several time-intervals in Tables 4 and 5 in the

appendix give additional evidence in favor of the non-stability of the FX rates distribution¹⁹.

From Table 3, one can also verify the invariance of the tail index under aggregation, except for the longest intervals, mainly the 6 hour interval where the small number of data becomes a problem in getting significant estimates of α . The smaller number of data for large intervals forces the estimation algorithm to use a larger fraction of this data, closer to the center of the distribution. Thus the empirically measured tail properties become distorted by properties of the center of the distribution which, for $\alpha > 2$ and under aggregation, approaches the normal distribution (with $\alpha = \infty$) as a consequence of the central limit theorem.

4.1.2 Fact 2: Finite variance. Table 3 and Tables 4 and 5 in the appendix suggest that the variance exists for all currencies.

4.1.3 Fact 3: Symmetric distribution. It appears from Table 3 and Tables 4 and 5 in the appendix that the third moment exists and the distribution is symmetric for all FX rates except, possibly, the EMS cross-rates.

4.1.4 Fact 4: Decreasing leptokurticity. However, it appears from these tables that the fourth moment is not finite. Indeed, the larger the number of observations, the larger the kurtosis will be. At frequencies higher than 10 minutes, there seems to be some contradiction between the work of Goodhard and Figliuoli (1991) which claims that the leptokurticity starts to decrease at these frequencies, and the paper of Bollerslev and Domowitz (1993) which gives some evidence of a still increasing leptokurticity. One can show, however, that both results hold depending on whether one uses the linear interpolation method or the previous tick (see Definition 1 and below) to obtain price values at fixed time intervals at such frequencies. This is an example of the difficulty of making reliable analyses of quoted prices at frequencies higher than 10 minutes.

The absence of the fourth moment explains why the absolute value of the price changes has been found to be the best definition of the volatility, that is, the one which exhibits the strongest structure. Indeed, since the fourth moment of the distribution comes into the computation of the autocorrelation function of the variance, its structure will depend on the number of data points used.

4.2. Price formation process

The following three facts pertain to the short term (less than 10 minutes) behavior of the foreign exchange intra-daily price changes. It should be stressed that these facts characterise FXFX quotes as opposed to real transaction prices. They highlight the need for the definition of an effective price and the difficulties inherent to tick-by-tick analysis.

4.2.1 Fact 5: Negative first-order autocorrelation of the returns. Goodhart (1989) and Goodhart and Figliuoli (1991) first reported the existence of negative first-order autocorrelation of the price changes at the highest frequencies, which disappear once the price formation process is over. Goodhart (1989) also demonstrated that this negative auto-correlation is not affected by the presence (or absence) of major news announcements. Finally, Goodhart and Figliuoli (1992) showed that the resulting oscillations of the prices are not caused by bouncing prices between different geographical areas with different information sets. Note that this negative first-order auto-correlation of FXX quotes is in contrast with the absence of such auto-correlation of real transaction prices, at least for the very small data sample studied in Goodhart et al. (1995).

A first explanation of this fact may be divergent opinions among traders. The conventional assumption that the FX market is composed of homogeneous traders who would share the same views about the effect of news, so that no correlation of the prices would be observed – or at most, a positive auto-correlation. However, traders have diverging opinions about the impact of news on the direction of prices. A second – and complementary – explanation for this negative auto-correlation is the tendency of market makers to skew the spread in a particular direction when they have order imbalances (Bollerslev and Domowitz 1993; Flood 1994). A third explanation is that even without order imbalances or diverging opinions on the price, certain banks systematically publish higher bid/ask spreads than other. This could also cause the ask (bid) prices to bounce back and forth between banks (Bollerslev and Melvin 1994).

4.2.2 Fact 6: Discreteness of spread. Quoted spreads are discretely distributed with the major peak being 10 basis points followed by the 5, 15 and 7 peaks. All together those peaks account for more than 97% of the distribution (Bollerslev and Melvin, 1994). These conventional spreads have also evolved over the years, depending on the markets, 10 rather than 20 basis points being the most quoted spread in recent years for example as the size of the price has become smaller (Müller and Sgier 1992).

As explained in Definition 5 above, spreads mainly depend on the cost structure of the bank that is the market maker and on the habits of the market. As shown in Goodhart and Curcio (1991), individual banks usually quote two or three different spreads. When a market maker wants to push the price in a particular direction, he will tend to skew the spread in that direction but will use a spread of conventional size, generally 5 or 7; when he only wishes to trade or is uncertain about the direction the price should take, he will quote larger spreads with conventional values such as 10 or 15. Because different banks have different conventions, the distribution of spreads has 4 or 5 peaks instead of 2 or 3.

Note, once again, that the quoted spreads do not exactly reflect the transaction spread, which are usually smaller except for highly volatile periods (Goodhart et al., 1995a). Another difference is the continuous distribution (between 0 and 20 basis points) of transaction spreads.

Consistent with theory (Admati and Pfleiderer 1988; Subrahmanyam 1991) market makers will cover themselves by conventionally larger spreads in periods of higher risk such as the release of important news (Goodhart 1989), the closing or opening of markets (Bollerslev and Domowitz 1993) and lunch breaks (Müller et al. 1990). More generally, the amplitude of the spread is inversely related to expected market activity as measured by tick frequency or mean hourly volatility (Müller et al. 1990). The amplitude of the spread is directly related to the (instantaneous) volatility, which also measures the risk (Bollerslev and Domowitz 1993).

4.2.3 Fact 7: Extremely short term triangular arbitrage. The building-up process of the prices in the very short run is also reflected in the significant predicting power of the USD/DEM relative to the other currencies (Goodhart and Figliuoli 1991). A short delay is indeed needed before traders in smaller currencies adjust themselves to the patterns of the two leading currencies. Eben (1994) also finds evidence of triangular arbitrage opportunities at very high frequencies arising when very short-term trend reversals between two USD-rates are not yet reflected in the quoted cross-rates. Although the detection of triangular arbitrage opportunities is rather easy and quick with a unique unit of account or vehicle currency, it takes more time when the rates between two vehicles (USD and DEM) change (Suvanto 1993). Finally, no study has yet shown whether such arbitrage opportunities are available with real transaction data.

4.3. Heterogeneous structure

The following facts show that the diversity of the market participants described in the first part is reflected in geographically and institutionally differentiated behaviors and in the presence of agents acting within different time-horizons. They also illustrate the information flows between these different market components. Note that these facts pertain to FX rates at frequencies lower than 10 to 15 minutes. As such, they are thus independent of the indicative nature of FXX quotes.

4.3.1 Fact 8: Seasonality. Although the FX market is virtually a global market, strong (deterministic) seasonal patterns corresponding to the hour of the day, the day of the week and the presence of the traders in the three major geographical trading zones can be observed. These seasonal patterns are found for all the five variables, the volatility (Bollerslev and Domowitz 1993; Dacorogna et al. 1993), the relative spread (Müller and Sgier 1992), the tick frequency (Goodhart and Demos 1990; Müller et al. 1990), the volatility ratio and the directional change frequency.

As can be seen on Figs. 1 and 2, trading activity picks up after midnight as the Tokyo and Sydney markets open with subsequent activity in Singapore




Fig. 1. Hourly intra-day and intra-week distribution of the absolute price change, the spread and the tick frequency. A sampling interval of $\Delta t = 1$ hour is chosen. The day is subdivided into 24 hours from 0:00 - 1:00 to 23:00 - 24:00 (GMT) and the week is subdivided into 168 hours from Monday 0:00 - 1:00 to Sunday 23:00 - 24:00 (GMT) with index i . Each observation of the analyzed variable is made in one of these hourly intervals and is assigned to the corresponding subsample with the correct index i . The sample pattern is independent of bank holidays and daylight saving time. The currency is the USD/DEM

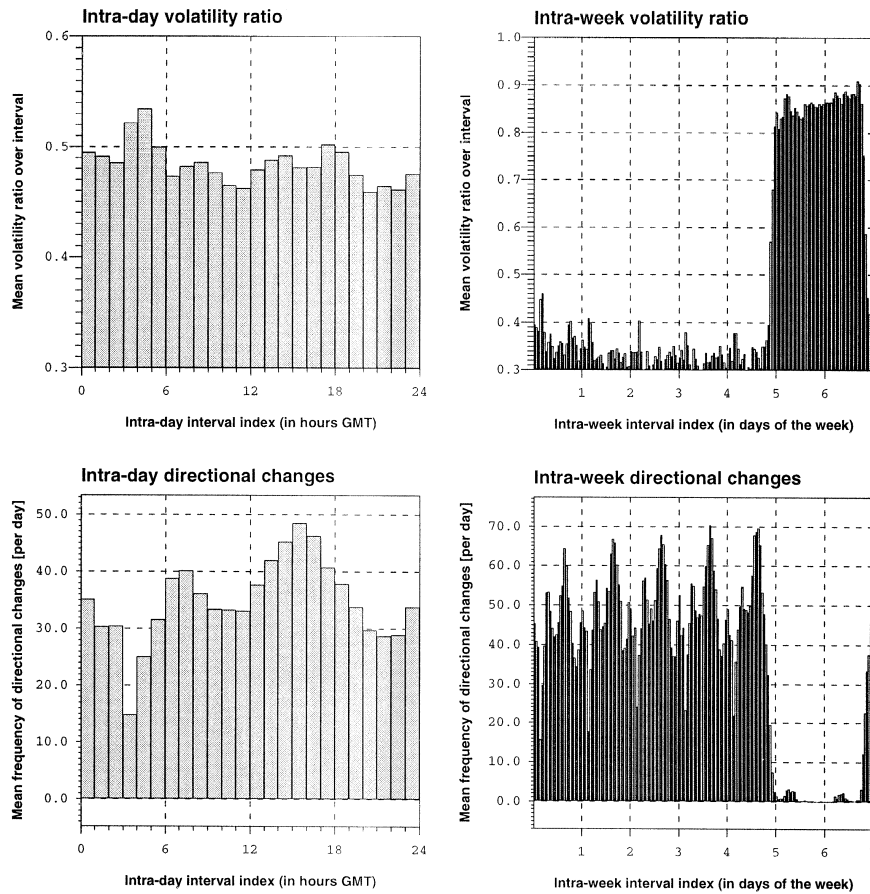


Fig. 2. Hourly intra-day and intra-week distribution of the volatility ratio and the directional change frequency. The number of subintervals (per hour) is 10. The threshold value for the directional change frequency is 0.0003. A sampling interval of $\Delta t = 1$ hour is chosen. The day is subdivided into 24 hours from 0:00 - 1:00 to 23:00 - 24:00 (GMT) and the week is subdivided into 168 hours from Monday 0:00 - 1:00 to Sunday 23:00 - 24:00 (GMT) with index i . Each observation of the analyzed variable is made in one of these hourly intervals and is assigned to the corresponding subsample with the correct index i . The sample pattern is independent of bank holidays and daylight saving time. The currency is the USD/DEM

and Hong Kong. The abrupt decline in arrivals at 4:00 GMT signals lunchtime in these markets. Market intensity remains strong in the afternoon Far Eastern trading session, and continues as Hong Kong and Singapore close and London and Frankfurt open. Some decline thereafter is observed until the opening in New York. Activity bounces back during the overlap of the New York and European markets, declining monotonically after New York closes and until the Far Eastern markets open again. Activity is regular over the five weekdays. With the exception of late Friday, which is already Saturday in the Australian and Asian markets, there is no real discernable difference across the weekdays. The average activity is low on weekends, but higher on Sundays than on Saturdays. That may be

due to the Middle Eastern markets, which can be open on Sundays and to early Monday morning activity in East Asia which coincides with Sunday nights. The seasonal patterns also appear in an autocorrelation study (see Fig. 3a).

Autocorrelation coefficients are indeed significantly higher for time lags that are integer multiples of the seasonal period than for other lags.

Although these seasonal patterns have similarities among all currencies, they can also be specific to a particular market. For CHF, for instance, there is a general spread increase with the high activity in America, while the average JPY spreads decrease at the same time. This indicates that American traders are less interested in CHF and more in JPY.

As shown in Andersen and Bollerslev (1994) and Guillaume et al. (1994a), the presence of these seasonal patterns introduces strong bias in the computation of simple statistics or the estimation of statistical processes with intra-daily data. This can lead to spurious results when this seasonality is not taken care of, as is the case in many empirical studies on intra-daily FX rates. There are however several ways to treat this seasonality. The first and most straightforward one is to introduce seasonal dummies as in Baillie and Bollerslev (1989). Another framework introduced in Andersen and Bollerslev (1994) is the use of a flexible Fourier framework to model the frequencies corresponding to the different seasonal peaks. A third possibility presented in Dacorogna et al. (1993) is the use of a different time scale, called Theta-time (ϑ -time). The ϑ -time scale expands day-times having high mean volatilities and contracts daytimes having low volatilities and the weekends having very low volatilities. As can be seen from Fig. 3b, the deterministic seasonal patterns almost vanish with this new time scale.

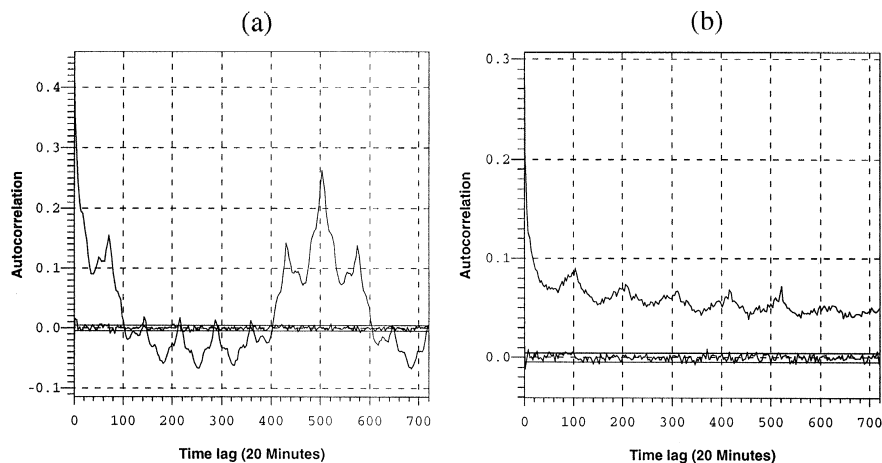


Fig. 3. Autocorrelation function. Autocorrelation function for the absolute price changes in physical time with seasonal peaks (a) and in Theta-time (ϑ -time) with long memory effects (b) for the GBP/USD. The horizontal curve and the two horizontal lines around zero are, respectively, the autocorrelation function of the price changes and the 95% confidence interval. The ϑ -time scale basically expands periods of high activity and contracts periods of low volatility. The weekends and business holidays are thus virtually omitted by this ϑ -scale (Dacorogna et al. 1993)

4.3.2 Fact 9: Short and long term memory. The high short-term autocorrelation of the volatility and its clustering in periods of high volatility and low volatility are well-known since the work of Mandelbrot (1963). This motivated the introduction of the Autoregressive Conditional Heteroskedastic (ARCH) model by Engle (1982) and its generalized version (GARCH) by Bollerslev (1986), which have had much success (see Bollerslev et al. 1992). More recently, it was found that the volatility also exhibits a very long-term memory characterized by an hyperbolic decline rather than the exponential decline of ARCH and GARCH processes. This led to another generalization of the GARCH model by Baillie et al. (1993), namely the Fractionally Integrated GARCH (FIGARCH).

Intra-daily studies not only confirm the presence of these short- and long-term memories for the volatility (Dacorogna et al. 1993) (see Fig. 3b) but also for other variables such as the spread (Müller and Sgier 1992), the volatility ratio and the directional change frequency.

Furthermore, the analysis of intra-daily data gives some insights on the origin of these clusters. One possibility is indeed the clustering of the arrival of news with the markets adjusting perfectly and instantaneously to the news. This is, however, not the reality as will be pointed out in Fact 14. Another possibility lies in the learning process of traders with different priors who may take some hours of trading to resolve their expectational differences after the arrival of important news. This would result in volatility spillovers that could either extend to other geographical trading areas or be restricted to their own market. Taking analogies from meteorology, Engle et al. (1990, 1992) called country-specific shocks a “heat wave” and the transmission of news across markets a “meteor shower”. Due to their limited (daily) data set, Engle et al. (1990,1992) only found evidence of meteor showers. However, using hourly data, Baillie and Bollerslev (1989) conclude to the presence of both the heat wave and the meteor shower phenomena. Further evidence is also given in Fig. 3b: on the one side, the continuously high short-term autocorrelation is responsible for the meteor shower; on the other side, the peaks at time lags corresponding to 1, 2 and 3 business days are evidence of the heat wave phenomenon. In addition to the presence of country-specific news, this heat wave could also be explained by the learning process of traders within a market. As one trader in the European market puts it, “when we arrive in the morning, we first take the temperature of the markets by reading the news, analyzing the price levels and inquiring about the mood of other traders. To complete a strategy for the day, we integrate customer orders and technical resistance and support levels. Only then we start with trading and may close positions before lunchtime or before the announcement of important news.” The response of traders to several different autocorrelated news arrival processes corresponding to different time-horizons is also a potential explanation for the long-term memory of the markets (Haubrich and Lo 1992).

4.3.3 *Fact 10: Fractal structure.* Another very striking fact is the regular fractal structure of the FX rates in the sense of Mandelbrot (1983)²⁰. This is illustrated by the scaling laws reported for the volatility and the directional change frequency.

The scaling law for the volatility (Müller et al. 1990) relates the volatility over a time interval Δt (see Definition 3) to the size of this interval:

$$v(\Delta t, S; t_i) = \left(\frac{\Delta t}{\Delta T} \right)^{\frac{1}{E}} \tag{12}$$

where S is the sampling period and ΔT is a constant depending on the FX rate. If Δt is expressed in days, ΔT is in the order of magnitude of 10^4 for the main FX rates against the USD. The drift exponent $1/E$ is about 0.58 for the major FX rates, compared to the $1/E = 0.5$ implying a pure gaussian random walk model.




Fig. 4. Scaling law for the volatility. This figure gives the scaling law for the volatility for intervals extending from 10 minutes to two months. The currency is the USD/JPY

A similar law relates the number of directional changes to the size of the threshold r_c :

$$S \cdot d(S, r_c; t_i) = \left(\frac{r_c}{R} \right)^{\frac{1}{D}} \quad (13)$$

where S is the sampling period and R is a constant depending on the FX rate. The drift exponent $1/D$ is about -1.75 for the major FX rates, whereas a pure gaussian random walk model would imply $1/D = -2.0$. Both scaling laws were estimated by least squares.




Fig. 5. Scaling law for the directional change. This figure gives the scaling law for the number of directional changes (see Definition 7) starting from threshold values of the size of the spread to the higher threshold values taken in a geometric progression. The currency is the USD/JPY

The scaling laws expressed in equations 12 and 13 hold for all time series studied and for a wide variety of time intervals – from 10 minutes to 2 months – or threshold values as shown in Figs. 4 and 5. Even more interestingly, recall that the definition of the directional change frequency is symmetric to the definition of the volatility: the volatility measures the variability of the price changes as

a function of a fixed time interval whereas the directional change frequency measures the variability of the price changes as a function of fixed amplitude or threshold. Then, taking the inverse of the drift exponent in the scaling law for the directional change frequency, we should get a value similar to the value of the drift exponent in the scaling law for the volatility; this is indeed approximately the case after taking care of the negative slope.

Although no theoretical models yet provide an explanation, a highly tentative economic interpretation of this scaling law is that it represents a mix of risk profiles of agents trading at different time horizons. The average volatility on one horizon is indeed the maximum return a trader can expect to make on average at that horizon. Alternatively, the average number of directional changes for a particular threshold or return is the maximum number of profitable trades a trader can expect to make on average. As shown in Fig. 6a, this relationship between traders with different risk or time-horizon profiles is quite stable over the years, notwithstanding the tripling of the volume on the FX markets.

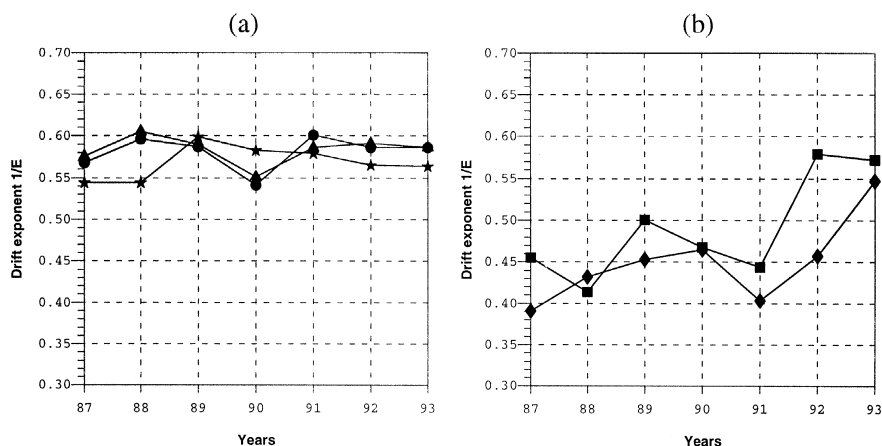


Fig. 6. Drift exponent as a function of time. Yearly estimation of the drift exponent of the scaling law (see Fact 10) is given for USD rates on the left (DEM (circle), FRF (triangle), JPY (star)) (a) and EMS rates against the DEM on the right (ITL (box) and FRF (diamond)) (b)

4.3.4 Fact 11: Institutional framework. Intra-daily analysis of FX rates within the European Monetary System (EMS) gives some insights on the distinct characteristics of the EMS system when the bands were still quite narrow. As illustrated in Fig. 6, it achieved a smaller average volatility, represented by a value of the drift exponent of the scaling law much smaller than the average for free-floating currencies. Though, the value of its drift exponent went up to the average value for free floating currencies when the ITL left the EMS in 1992 and when the bands were broadened in 1993 in the case of the FRF.

The earlier volatility was achieved however at the cost of a much larger probability of extreme events as indicated by the smaller value of the tail index of

EMS currencies in comparison with free-floating currencies (see Table 3 in Fact 1). This statistical characteristic demonstrates the low credibility of the former EMS' bands during our sample period (01.10.1992-31.12.1993) and shows that different institutional set-ups like the EMS can be distinguished by distinct and robust, model or process-independent, features²¹.

4.3.5 Fact 12: Conditional predictability. The co-existence of different types of traders might be why conditional forecasts are possible although the price changes are globally unpredictable (see Fact 15). Based on earlier work by LeBaron (1992), Müller et al. (1994) show that tomorrow's volatility on a very short time horizon (for example 1 hour) can be systematically predicted by the volatility on a longer time-horizon (for example 1 day). Besides, Guillaume (1996) also shows that the more trend-following the behaviour of traders was yesterday, the higher the volatility will be today; conversely, tomorrow's trend-following behaviour is negatively predicted by today's height of the volatility. These facts seem to indicate that the impact of important news first affects middle-term traders (over 1 day) and then propagate itself to short-term (intra-day) traders, who have to absorb the shock by a more intense activity. On the other hand, a more intense activity between intra-daily dealers will not cause trend in the relatively longer run to appear. This is in line with the fact that most of the intra-daily volume is made on the basis of "hot potatoes" exchanged between different intra-daily dealers (Lyons 1994). These findings illustrate once again the change of dynamics for different types of volatility and the importance of information flows between clearly distinct types of traders.

4.3.6 Fact 13: Positive impact of official interventions. One type of trader that is of special interest are the central banks as the time and the size of their interventions can be measured on an intra-daily basis. Central banks may operate either directly through officially announced interventions, indirectly through unannounced interventions or through big banks which represent a larger share of the market. Official interventions operate essentially as signals given to the markets and are therefore difficult to measure (see Edison (1993) for a review of the literature on central bank interventions).

Some evidence is given in Goodhart and Hesse (1993) that in the long run, official interventions would affect FX rates in the desired direction although they may result in short term losses for the central bankers. One could, however, easily extend the analysis to any other long-term trader - that is, a trader who can afford a large open position - provided he has some impact on the market through his reputation even if he does not have a large share of the market. For example, this may be the case of some hedge funds. Peiers (1994) shows the positive impact of indirect interventions of the Bundesbank through the biggest player on the DEM/USD market, namely the Deutsche Bank.

4.3.7 *Fact 14: Mixed effect of news.* News is a very broad concept covering a phone call of a customer who wants to make a very large FX transaction (due to, say, inventory imbalances), a conversation with a colleague, price forecasts and histories when used in technical analysis programs or the economic forecasts of the research department of a bank, general economic and political news and major economic news announcements. News is therefore very difficult to quantify.

A first attempt to quantify news taken in Goodhart (1989) was to look at the ‘news’ pages of Reuters. General economic and political news are displayed on the AAMM page. Goodhart (1989) found out that unimportant news does not have a significant effect on the behaviour of the foreign exchange rates. Distinct and relatively large price movements unrelated to any news are indeed apparent. The price formation process seems to be independent of the presence or absence of news.




Fig. 7. Intraday distribution per 15 minutes interval. Intraday distribution per 15 minutes interval of the second difference of the mean absolute price changes for the USD/DEM ($h_i = \exp[\ln|r_i| - \frac{1}{2} \ln|r_{i-1}r_{i+1}|]$). Sudden peaks are in darkened

In contrast, major economic news announcements such as trade, unemployment, budget deficit or gross domestic product growth have significant impacts (Goodhart 1989). Economic news announcements along with the market expectations and the effect of the previous announcement are displayed in Reuters' FXNB page. Effective news – that is, the difference between the markets expectation and the actual figure that is released – increase the volatility as the dispersion of traders' views on the impact of the effective news widens. The three right peaks in Fig. 7 show the clearcut effect of this news release in New York and Japan. The two peaks for the US reflect the change of time with the Daylight Saving Time, which does not exist in Japan. The first two peaks on the left correspond to the beginning of the Japanese trading session and of the Japanese lunch respectively. Goodhart et al. (1993) further shows that major economic news announcement such as release of US trade figures or changes in the UK base interest rate have a significant impact on the price changes process. This effect is, however, temporary (3 to 4 days) as markets eventually incorporate the effects of the news. Moreover the direction of the effect on the level of the price is difficult to predict. This can be explained by the highly nonlinear dynamics of the FX rates (Guillaume et al. 1996).

An alternative way to quantify the impact of news derives from the mixture of distribution hypothesis (see Clark 1973; Tauchen and Pitts 1983; Andersen 1992). In this framework, the clustering of the volatility results from the clustering of the news arrival process. Since the news arrival process is an unobserved variable, proxies for the market activity such as the volume of trade are used. Volume as such is however not available in the FX markets (see Sect. 2). Moreover, as shown in Jones et al. (1991), volume can be rather noisy. Therefore, empirical studies in the FX intra-daily markets use the tick frequency and/or the spread as proxies for the level of activity. Although a certain correlation between these variables and the volatility is obvious from the simple inspection of Fig. 1 in Fact 5 above, severe limitations harm the use of these variable as noted in Sect. 2. Moreover Davé (1993) shows that tick frequency can only be a good approximation of the volume when markets are analyzed as separate geographical entities – thus, when there is no overlap between markets and the data are not disaggregated by individual bank subsidiary. Goodhart (1989) also shows that tick frequency does not specifically rise when news are released. Anyhow, empirical evidence in favor of this mixture of distribution hypothesis is only partial (Bollerslev and Domowitz 1993; Demos and Goodhart 1992).

4.3.8 Fact 15: Highly complex system. As implied by the presence of a unit root (Goodhart and Figliuoli 1992), FX price changes appear to be highly unpredictable. Though, at the same time, their dynamics is highly nonlinear (Guillaume et al. 1996) and, at least partly, endogenous as highlighted by Fact 14. This endogeneity of the price movements results from the information flow between agents trading within different geographical markets and time-horizons.

Since this endogeneity, together with the nonlinearity and the unpredictability, is the characteristic of a chaotic system, an interesting property to investigate is the degree of complexity of this chaotic system. A low-degree of complexity would indicate that the FX rates dynamics can mainly be described by an endogenous system with some small stochastic perturbations. Rather than the news itself, it is the complex nonlinear interaction and learning process of traders that would be responsible for the large and unpredictable movements of FX rates (DeGrauwe and Dewachter 1993). Using a very large number of intra-daily data over a sufficiently long period, Guillaume (1995) found however that FX rates are a highly complex system.

Together with Fact 14, this fact implies that the dynamics of the FX markets cannot be simply explained by either the impact of one news process or the endogenous learning process of traders. More systematic relationships between the FX rates and the fundamentals or the news can probably only be detected by first trying to isolate the different components of the markets corresponding to different types of traders.

5. Conclusion

Looking through the microscope at the intra-daily FX market, we have access to much more information than at lower frequencies such as daily or weekly. Not only is the number of data points 100 to 1000 times larger, other types of information such as news screens or new quantities are now available. This information set is highly interesting as it is the one used by intra-daily traders who make up by far the largest share of the market (Bank of International Settlement 1993). Its analysis also offers answers to several old debates and poses new challenges.

A first issue is the definition of an effective price at the extremely high frequency to replace transaction prices. Indeed, in the absence of widely available transaction prices, an effective price algorithm could be devised on the basis of the characteristics of the price formation process and the market structure organisation.

A second and somewhat related issue is the definition of the relevant time-scale. Contrary to traditionally available data, ticks are irregular spaced in physical time. Moreover, the heterogeneity of the markets causes inconsistencies in the estimation of models in physical time. Therefore, several alternative time-scales were suggested.

A third issue is the definition of the concept of risk. Risk is traditionally defined by the volatility as measured by the variance and the co-variance. However, the absolute value of the returns is a more accurate definition of the volatility due to the non-existence of the fourth moment of the distribution. The tail index is another measure of risk which can be more accurate than the variance and co-variance. The analysis of EMS FX rates indeed showed that a reduced volatility does not necessarily imply a lower risk as measured by the tail index.

A fourth challenge is the definition of market efficiency and the related issue of the learning process of traders. The need for a dynamic definition of the concept of market efficiency is stressed by the mixed impact of news on the price changes on the one hand, and the spillover effects both accross markets and accross time-horizons on the other hand. This dynamic view of market efficiency does not preclude the possibility of conditional forecasting of the price changes as it explicitly takes into account the heterogeneity of the markets. Rather than the instantaneous adjustment of the price to news, market efficiency can now be defined as the smoothness of this adjustment process. In this perspective, appropriate measures of market efficiency could be the size of the spread or the tail index.

With respect to the heterogeneous behaviour of the traders, a fifth challenge is the definition of the right modelling approach to the markets. The short-comings of both the macro-economic and the time-series methodologies adopted so far have been highlighted. A first alternative may reside in the analysis of the sub-components of the markets rather than the market as a whole by using, for example, other time-scales corresponding to the time-horizon of a specific category of traders. A second alternative is the analysis of a particular category of traders whose specific time-horizon and objective function can be identified, like individual traders who would communicate the record of their trading. A third alternative would be to broaden the conceptual framework of the theoretical literature on the micro-structure of the markets to include traders with different time-horizons or geographical locations for example. Without taking this heterogeneity of the markets into account, it may prove extremely difficult if not impossible to find relationships between the price changes and the fundamentals or more generally the news arrival processes. For example, even the simple geographical dispersion of the different markets proves to be of prime importance for the consistent estimation of simple statistics or statistical models.

Finally, some light was shed on the old debate on the distribution of FX returns which was shown to belong to the class of fat-tailed non-stable distributions. Furthermore, the existence of the three first moments of the distribution for free-floating currencies was confirmed.

6. Appendix

6.1. Appendix 1

Table 4. Price change distribution. This table gives four first moments of the price change distribution at different time intervals for the major currencies against the USD

Rate	Time interval	Mean	Variance	Skewness	Kurtosis
USD/DEM	10 minutes	$-2.73 \cdot 10^{-7}$	$2.62 \cdot 10^{-7}$	0.17	35.10
	1 hour	$-1.63 \cdot 10^{-6}$	$1.45 \cdot 10^{-6}$	0.26	23.55
	6 hours	$-9.84 \cdot 10^{-6}$	$9.20 \cdot 10^{-6}$	0.24	9.44
	24 hours	$-4.00 \cdot 10^{-5}$	$3.81 \cdot 10^{-5}$	0.08	3.33
	1 week	$-2.97 \cdot 10^{-4}$	$2.64 \cdot 10^{-4}$	0.18	0.71
USD/JPY	10 minutes	$-9.42 \cdot 10^{-7}$	$2.27 \cdot 10^{-7}$	-0.18	26.40
	1 hour	$-5.67 \cdot 10^{-6}$	$1.27 \cdot 10^{-6}$	-0.09	25.16
	6 hours	$-3.40 \cdot 10^{-5}$	$7.63 \cdot 10^{-6}$	-0.05	11.65
	24 hours	$-1.37 \cdot 10^{-4}$	$3.07 \cdot 10^{-5}$	-0.15	4.81
	1 week	$-9.61 \cdot 10^{-4}$	$2.27 \cdot 10^{-4}$	-0.27	1.30
GBP/USD	10 minutes	$-6.91 \cdot 10^{-9}$	$2.38 \cdot 10^{-7}$	0.02	27.46
	1 hour	$7.61 \cdot 10^{-7}$	$1.40 \cdot 10^{-6}$	-0.23	21.53
	6 hours	$4.63 \cdot 10^{-6}$	$8.85 \cdot 10^{-6}$	-0.34	10.09
	24 hours	$1.72 \cdot 10^{-5}$	$3.60 \cdot 10^{-5}$	-0.26	4.41
	1 week	$6.99 \cdot 10^{-5}$	$2.72 \cdot 10^{-4}$	-0.66	2.77
USD/CHF	10 minutes	$-2.28 \cdot 10^{-7}$	$3.07 \cdot 10^{-7}$	-0.04	23.85
	1 hour	$-1.37 \cdot 10^{-6}$	$1.75 \cdot 10^{-6}$	0.05	18.28
	6 hours	$-8.23 \cdot 10^{-6}$	$1.11 \cdot 10^{-5}$	0.05	7.73
	24 hours	$-3.38 \cdot 10^{-5}$	$4.51 \cdot 10^{-5}$	-0.04	2.81
	1 week	$-2.58 \cdot 10^{-4}$	$3.16 \cdot 10^{-4}$	0.09	0.34
USD/FRF	10 minutes	$-1.98 \cdot 10^{-7}$	$2.08 \cdot 10^{-7}$	0.35	43.31
	1 hour	$-1.18 \cdot 10^{-6}$	$1.28 \cdot 10^{-6}$	0.47	28.35
	6 hours	$-7.13 \cdot 10^{-6}$	$8.29 \cdot 10^{-6}$	0.23	9.69
	24 hours	$-2.91 \cdot 10^{-5}$	$3.40 \cdot 10^{-5}$	0.06	3.22
	1 week	$-2.32 \cdot 10^{-4}$	$2.44 \cdot 10^{-4}$	0.16	0.88

Table 5. Price change distribution. This table gives four first moments of the price change distribution at different time intervals for the major currencies against the DEM

Rate	Time interval	Mean	Variance	Skewness	Kurtosis
DEM/FRF	10 minutes	$9.84 \cdot 10^{-8}$	$1.91 \cdot 10^{-8}$	0.54	86.29
	1 hour	$5.89 \cdot 10^{-7}$	$1.14 \cdot 10^{-7}$	0.79	69.70
	6 hours	$3.53 \cdot 10^{-6}$	$6.53 \cdot 10^{-7}$	1.41	36.87
	24 hours	$1.07 \cdot 10^{-5}$	$2.84 \cdot 10^{-6}$	1.15	24.26
	1 week	$8.94 \cdot 10^{-5}$	$1.93 \cdot 10^{-6}$	1.92	3.95
DEM/NLG	10 minutes	$-5.19 \cdot 10^{-8}$	$1.42 \cdot 10^{-9}$	-5.68	9640.85
	1 hour	$-3.11 \cdot 10^{-7}$	$7.54 \cdot 10^{-9}$	2.76	4248.12
	6 hours	$-1.86 \cdot 10^{-6}$	$2.48 \cdot 10^{-8}$	0.74	124.35
	24 hours	$-7.80 \cdot 10^{-6}$	$9.66 \cdot 10^{-8}$	-0.30	30.02
	1 week	$-4.57 \cdot 10^{-5}$	$6.63 \cdot 10^{-7}$	0.03	0.06
DEM/ITL	10 minutes	$1.07 \cdot 10^{-6}$	$1.75 \cdot 10^{-7}$	0.86	64.03
	1 hour	$6.46 \cdot 10^{-6}$	$1.24 \cdot 10^{-6}$	1.83	89.92
	6 hours	$3.88 \cdot 10^{-5}$	$7.16 \cdot 10^{-6}$	1.03	37.26
	24 hours	$1.18 \cdot 10^{-4}$	$2.53 \cdot 10^{-5}$	-0.51	13.08
	1 week	$9.42 \cdot 10^{-4}$	$1.37 \cdot 10^{-4}$	-0.25	0.17
GBP/DEM	10 minutes	$4.53 \cdot 10^{-7}$	$9.86 \cdot 10^{-8}$	-0.32	25.97
	1 hour	$2.69 \cdot 10^{-6}$	$7.12 \cdot 10^{-7}$	-0.34	16.90
	6 hours	$1.56 \cdot 10^{-5}$	$4.62 \cdot 10^{-6}$	-0.02	7.48
	24 hours	$7.04 \cdot 10^{-5}$	$1.79 \cdot 10^{-5}$	0.27	3.15
	1 week	$1.17 \cdot 10^{-4}$	$1.29 \cdot 10^{-4}$	0.07	0.59
DEM/JPY	10 minutes	$-3.39 \cdot 10^{-6}$	$2.21 \cdot 10^{-7}$	-0.09	12.35
	1 hour	$-2.03 \cdot 10^{-5}$	$1.46 \cdot 10^{-6}$	-0.03	88.58
	6 hours	$-1.21 \cdot 10^{-4}$	$9.12 \cdot 10^{-6}$	-0.04	6.57
	24 hours	$-4.85 \cdot 10^{-4}$	$3.56 \cdot 10^{-5}$	0.12	2.52
	1 week	$-3.15 \cdot 10^{-3}$	$2.67 \cdot 10^{-4}$	-0.07	0.03

6.2. Appendix 2

A more detailed definition of the *directional change frequency* at time t_i , $d(t_i)$, is as

$$d(t_i) \equiv d(\Delta t, n, r_c; t_i) \equiv \frac{1}{n \Delta t} N(\{k \mid m_k \neq m_{k-1}, 1 < k \leq n\}) \quad (14)$$

with the recursive variables m_k , which indicate the mode – upwards or downwards – of the current trend, and min_k (max_k), which indicate the minimum (maximum) value used as a reference to compute potential changes of mode. Formally, we have:




Fig. 8. Schematic representation of the directional change algorithm. This figure described how the algorithm for the directional change frequency (see Definition 7) is implemented. The mode is only initialised to the upward mode at time t_2 since the first change at time t_1 is not significant as defined by r_c . At time t_3 , the downward movement is not taken into account as it is smaller than the significance level r_c . At time t_4 , we are thus still in the upward mode with the maximum value being adjusted according to the algorithm. At time t_5 , the mode is reversed as the move downward is significant. At the same time, the minimum value is adjusted up to min_5 . At time t_6 , the minimum is again adjusted as we are still in a down mode. At time t_7 , the mode is reversed to the up mode and the maximum is adjusted as the change upwards is significant. There are therefore two directional changes in total

$$m_k = \begin{cases} 0 & , k = 0 \\ \left\{ \begin{array}{l} +1 \quad \text{if } ((m_{k-1} \neq +1) \wedge (x(t_{i-n+k}) - min_{k-1} > r_c)) \\ -1 \quad \text{if } ((m_{k-1} \neq -1) \wedge (max_{k-1} - x(t_{i-n+k}) > r_c)) \\ m_{k-1} \quad \text{otherwise} \end{array} \right\} & , k = 1, \dots, n \end{cases}$$

$$max_k = \begin{cases} x(t_{i-n+k}) & \text{if } (k = 0) \vee [(m_k = 1) \wedge (m_{k-1} \neq 1)] \\ \max(max_{k-1}, x(t_{i-n+k})) & \text{otherwise} \end{cases}$$

$$min_k = \begin{cases} x(t_{i-n+k}) & \text{if } (k = 0) \vee [(m_k = -1) \wedge (m_{k-1} \neq -1)] \\ \min(min_{k-1}, x(t_{i-n+k})) & \text{otherwise} \end{cases}$$

where the computation sequence is $m_0, max_0, min_0, m_1, max_1, min_1, \dots$ and where $N(\{k\})$ is the counting function, $n\Delta t$ the sampling period on which the counting is performed and r_c is a threshold value. The directional change frequency, $d(t_i)$, is simply the frequency of significant mode (m_k) changes with respect to the latest extremum value (max_k or min_k) and a constant threshold value r_c .

The threshold, r_c , should be related to the time frequency ($1/\Delta t$) at which the price change is taken and to the mean absolute value of the price change for this time interval Δt . On the one hand, the value of r_c should be large enough not to measure “noise” and, on the other hand, r_c should also reflect the typical sensitivity level of a trader operating at the corresponding time-horizon Δt . The scaling law described in Fact 7 may represent a possibility of linking r_c to Δt .

Endnotes

¹ For surveys on the FX market at the daily or weekly frequencies, see, for example, the surveys of Mussa (1979), Hsieh (1988), Baillie and McMahon (1989), de Vries (1994).

² Holidays and business hours for the different markets can be found in Morgan (1994).

³ East Asia comprises Australia, Hong Kong, India, Indonesia, Japan, South Korea, Malaysia, New Zealand, Singapore. Europe comprises Austria, Bahrain, Belgium, Germany, Denmark, Finland, France, Great Britain, Greece, Ireland, Italy, Israel, Jordan, Kuwait, Luxembourg, Netherlands, Norway, Saudi Arabia, South Africa, Spain, Sweden, Switzerland, Turkey, United Arab Emirates. America comprises Argentina, Canada, Mexico, USA.

⁴ Standard abbreviations of the International Organization for Standardization (ISO, code 4217).

⁵ Net-net turnover; that is, adjusted for both local and cross-border double-counting.

⁶ The high leverage and unregulated aspects of hedging funds distinguish their investors from other institutional investors.

⁷ For example, Bank Negara of Malaysia was one of the most aggressive (short-term) speculators in the FX market for several years.

⁸ The spot market accounted for 47% of the FX market in 1992 (BIS 1993).

⁹ In the United States and the United Kingdom, the share of deals going through such systems in April 1992 was 32 and 28% respectively (BIS 1993).

¹⁰ Lyons (1993a,b, 1994) and Goodhart et al. (1994) could exceptionally get respectively one week and one day of such data.

¹¹ See Guillaume and Payne (1996) for a comparison between Reuters' FXFX system and its new RIC system.

¹² In their one day study of real transaction prices, Goodhart et al. (1994) found that although the actual spread is usually within the quoted spread, it could be larger in highly volatile periods.

¹³ A trader taking a forward position overnight will of course have to pay the interest on his position between the trade and the settlement as well as the spread on the interest rate.

¹⁴ The back office usually takes longer term positions than traders in the trading room who might not be allowed to take positions overnight.

¹⁵ Appendix 2 gives a more detailed definition of the directional change frequency and a schematic example of its implementation.

¹⁶ By trade, we mean entering *and* closing a position.

¹⁷ The tail index represents a measure of tail fatness and can be derived from the condition of regular variation at infinity. Let $F(t)$ be a distribution function. Then if

$$\lim_{t \rightarrow \infty} \frac{1 - F(tx)}{1 - F(t)} = x^\alpha \quad (15)$$

holds for some α with α and x positive, $F(t)$ is said to be regularly varying with tail index α .

¹⁸ This is an implication of the central limit law.

¹⁹ Simulations in McCulloch (1994) and Gielens et al. (1995) show that one cannot univoquely distinguish between a fat-tailed non-stable and a thin-tailed distribution only on the basis of low estimated values of the tail index. However, the estimations of the kurtosis point in favor of the non-stability of the FX rates distribution.

²⁰ The fractal structure of the FX rates does not correspond to a fractional brownian motion since the distribution is non stable. It is also more complex than the structure of a low-dimensional fractal attractor as illustrated by Fact 15.

²¹ See Svensson (1992) for a review of the literature on the modeling of Target Zones, and in particular, of the EMS.

References

1. Admati, A.R., Pfleiderer, P.: A theory of intraday patterns: volume and price variability. *Rev. Financial Studies* **1**, 3–40 (1988)
2. Andersen, T.G.: Return volatility and trading volume: an information flow interpretation of stochastic volatility. Kellogg Graduate School of Management, Northwestern University, working paper 170, pp. 1–35 (1992)
3. Andersen, T.G., Bollerslev, T.: Intraday seasonality and = volatility persistence in foreign exchange and equity markets. Kellogg Graduate School of Management, Northwestern University, working paper (1994)
4. Baillie, R.T., Bollerslev, T.: The message in daily exchange rates: a conditional-variance tale. *J. Business Econ. Statist.* **7**(3), 297–305 (1989)
5. Baillie, R.T., Bollerslev, T., Mikkelsen, H.-O.: Fractionally integrated generalized autoregressive conditional heteroskedasticity. Kellogg Graduate School of Management, Northwestern University, working paper 186, pp 1–24 (1993)
6. Baillie, R.T., McMahon, P.C.: *The foreign exchange market*. Cambridge: Cambridge University Press 1989
7. Bank for International Settlements: Central bank survey of foreign exchange market activity. Document from the Monetary and Economic Department, Basle, February 1993, pp. 1–42 (1993)
8. Bollerslev, T.: Generalized autoregressive conditional heteroskedasticity. *J. Econometrics* **31**, 307–327 (1986)
9. Bollerslev, T., Chou, R.Y., Kroner, K.F.: ARCH modeling in finance. *J. Econometrics* **52**, 5–59 (1992)
10. Bollerslev, T., Domowitz, I.: Trading patterns and prices in the interbank foreign exchange market. *J. Finance* **48**, 1421–1443 (1993)
11. Bollerslev, T., Melvin, M.: Bid-ask spreads and volatility in the foreign exchange market: an empirical analysis. *J. Intern. Econ.* **36**, 355–372 (1994)
12. Boothe, P., Glassman, D.: The statistical distribution of exchange rates, empirical evidence and economic implications. *J. Intern. Econ.* **22**, 297–319 (1987)
13. Calderon-Rossel, J.R., Ben-Horim, M.: The behavior of the foreign exchange rates, empirical evidence and economic implications. *J. Intern. Business Studies* **13**, 99–111 (1982)
14. Clark, P.K.: A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* **41**(1), 135–155 (1973)
15. Cox, J.C., Rubinstein, M.: *Option markets*. Englewood Cliffs, NJ: Prentice Hall 1985
16. Dacorogna, M.M., Müller, U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V.: A geographical model for the daily and weekly seasonal volatility in the FX market. *J. Intern. Money Finance* **12**(4), 413–438 (1993)
17. Dacorogna, M.M., Pictet, O.V., Müller, U.A., de Vries, C.G.: The distribution of extremal foreign exchange rate returns in extremely large data sets. Unpublished manuscript, Olsen & Associates, Zürich, 1994
18. Davé, R.D.: Statistical correlation of data frequency price change and spread results. Unpublished manuscript, Olsen & Associates, Zürich, 1993
19. de Haan, L.: Fighting the arch-enemy with mathematics. *Statist. Neerlandica* **44**, 45–68 (1990)
20. de Vries, C.G.: Stylized facts of nominal exchange rate returns. In: Ploeg, F. van der (ed.): *Handbook of International Macroeconomics*. Blackwell, Oxford, pp. 348–389 (1994)
21. de Grauwe, P., Dewachter, H.: A chaotic model of the exchange rate: the role of fundamentalists and chartists. *Open Econ. Rev.* **4**, 351–379 (1993)
22. Demos, A.A., Goodhart, C.A.E.: The interaction between the frequency of market quotations, spread, and volatility in the foreign exchange market. LSE Financial Markets Group Discussion Paper **152**, 1–38 (1992)
23. Dunis, C., Feeny, M.: *Exchange rate forecasting*. Cambridge: Woodhead-Faulkner 1989
24. Eben, K.: Arbitrage alerts and changing interrelations between fx-rates. Unpublished manuscript, Institute of Computer Science, Academy of Sciences of the Czech Republic, 1994
25. Edison, H.: The effectiveness of central bank intervention: a survey of the Post-1982 literature. *Essays in International Finance*. Princeton University, 1993
26. Engle, R.F.: Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica* **50**, 987–1008 (1982)

27. Engle, R.F., Ito, T., Lin, W.-L.: Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market. *Econometrica* **58**, 525–542 (1990)
28. Engle, R.F., Ito, T., Lin, W.-L.: Where does the meteor shower come from: the role of stochastic policy coordination. *J. Intern. Econ.* **32**, 221–240 (1992)
29. Feller, W.: An introduction to probability theory and its applications. Vol. II of Wiley Series in Probability and Mathematical Statistics, 2nd edn. New York: John Wiley 1971
30. Flood, M.D.: Market structure and inefficiency in the foreign exchange market. *J. Intern. Money Finance* **13**(2), 131–158 (1994)
31. Gielens, G., Straetmans, S., de Vries, C.G.: Fat tail distributions and local thin tail alternatives. Mimeo, Tinbergen Instituut Rotterdam, pp. 1–14, 1995
32. Goodhart, C., Ito, T., Payne, R.: One day in June, 1993: a study of the working of the Reuters dealing 2000-2 electronic foreign exchange trading system. Mimeo, London School of Economics, pp. 1–36, 1995
33. Goodhart, C.A.: ‘news’ and the foreign exchange market. Proceedings of the Manchester Statistical Society, pp. 1–79, 1989
34. Goodhart, C.A., Demos, A.: Reuters screen images of the foreign exchange market: the deutschemark/dollar spot rate. *J. Intern. Securities Markets* **4**, 333–348 (1990)
35. Goodhart, C.A., Figliuoli, L.: The geographical location of the foreign exchange market: a test of an “islands” hypothesis. *J. Intern. Comp. Econ.* **1**, 13–27 (1992)
36. Goodhart, C.A.E., Figliuoli, L.: Every minute counts in financial markets. *J. Intern. Money Finance* **10**, 23–52 (1991)
37. Goodhart, C.A.E., Hall, S.G., Henry, S.G.B., Pesaran, B.: News effects in a high frequency model of the sterling-dollar exchange rate. *J. Appl. Econometrics* **8**, 1–13 (1993)
38. Goodhart, C.A.E., Hesse, T.: Central Bank Forex intervention assessed in continuous time. *J. Intern. Money Finance* **12**(4), 367–389 (1993)
39. Goodhart, C.C., Curcio, R.: The clustering of bid/ask prices and the spread in the foreign exchange market. LSE Financial Market Group Discussion Paper Series **110**, 1–15 (1991)
40. Granger, C., Ding, Z.: Some properties of absolute return: An alternative measure of risk. Presented at the conference on ‘Financial markets dynamics and forecasting’, organized by the Groupe Caisse des Depots and G.R.E.Q.E., pp. 1–30, 1993
41. Guillaume, D.: A low-dimensional fractal attractor in the foreign exchange markets? In: Trippi, R. (ed.): Chaos and nonlinear dynamics in the financial markets. Irwin Concorde, Ontario, pp. 269–294, 1995
42. Guillaume, D.: Chaos, randomness and order in the foreign exchange markets: Essays on the modelling of the markets. Unpublished Ph.D. thesis, University of Leuven, pp. 1–171, 1996
43. Guillaume, D., Payne, R.: Geographical information flows in the foreign exchange markets. Unpublished manuscript, Financial Markets Group, London School of Economics, 1996
44. Guillaume, D.M., Dacorogna, M.M., Pictet, O.V.: On the intra-daily performance of garch processes. Unpublished manuscript, Financial Markets Group, London School of Economics, 1994
45. Guillaume, D.M., Pictet, O.V., Müller, U.A., Dacorogna, M.M.: Unveiling nonlinearities through time scale transformations. Unpublished manuscript, Financial Markets Group, London School of Economics, 1996
46. Haubrich, J.G., Lo, A.W.: The sources and nature of long-term dependence in the business cycle. Unpublished manuscript, Department of Finance, M.I.T., 1992
47. Hsieh, D.A.: The statistical properties of daily foreign exchange rates: 1974–1982. *J. Intern. Econ.* **24**, 129–145 (1988)
48. International Monetary Fund: International capital markets. Part I. Exchange rate management and international capital flows, World Economic and Financial Surveys, Washington, April 1993, pp. 1–79
49. Jones, C.M., Kaul, G., Lipson, M.L.: Transactions, volumes and volatility. Mimeo, University of Michigan, School of Business Administration, 1991
50. Morgan, J.P.: World holiday and time guide. Morgan Kaufmann, Hove, UK
51. Levich, R.M., Thomas, L.R.III: The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach. *J. Intern. Money Finance* **12**(5), 451–474 (1993)
52. Lo, A.W., MacKinlay, A.C.: Stock market prices do not follow random walks: evidence from a simple specification test. *Rev. Financial Studies* **1**, 41–66 (1988)
53. Loretan, M., Phillips, P.C.B.: Testing the covariance stationarity of heavy-tailed time series. *J.*

- Empirical Finance **1**(2), 211–248 (1994)
54. Lyons, R.K.: Optimal transparency in a dealership market with an application to foreign exchange. NBER Working Paper Series no. 4467, pp. 1–32, 1993a
 55. Lyons, R.K.: Test of microstructural hypotheses in the foreign exchange market. NBER Working Paper Series no. 4471, pp. 1–32, 1993b
 56. Lyons, R.K.: Foreign exchange volume: sound and fury signifying nothing? Mimeo, U.C. Berkeley, 1994
 57. Mandelbrot, B.B.: The variation of certain speculative prices. *J. Business* **36**, 394–419 (1963)
 58. Mandelbrot, B.B.: *The fractal geometry of nature*. New York: W.H. Freeman 1983
 59. McCulloch, J.H.: Measuring tail thickness in order to estimate the stable index α : a critique. Mimeo, Ohio State University, pp. 1–34, 1994
 60. McFarland, J.W., Petit, R.R., Sung, S.K.: The distribution of foreign exchange price changes: trading day effects and risk measurement. *J. Finance* **37**(3), 693–715 (1982)
 61. Meyers, T.A.: *Technical Analysis Course*. Chicago: Probus Publishing Company 1989
 62. Müller, U.A.: Statistics of variables observed over overlapping intervals. Unpublished manuscript, Olsen & Associates, Zürich, 1993
 63. Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Pictet, O.V., von Weizsäcker, J.E.: Volatilities and trends of different time resolutions – analyzing the dynamics of market components. *J. Empirical Finance* (Forthcoming, 1997)
 64. Müller, U.A., Dacorogna, M.M., Davé, R.D., Pictet, O.V., Olsen, R.B., Ward, J.R.: Fractals and intrinsic time – a challenge to econometricians. Unpublished manuscript, Olsen & Associates, Zürich, 1993
 65. Müller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M., Morgeneegg, C.: Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *J. Banking Finance* **14**, 1189–1208 (1990)
 66. Müller, U.A., Sgier, R.G.: Statistical analysis of intra-day bid-ask spreads in the foreign exchange market. Unpublished manuscript, Olsen & Associates, Zürich, 1992
 67. Mussa, M.: Empirical regularities in the behavior of exchange rates and theories of the foreign exchange market. *Carnegie-Rochester Series on Public Policy*, pp. 9–57, 1979
 68. Neftci, S.N.: Naive trading rules in financial markets and Wiener-Kolmogorov prediction theory: a study of “technical analysis”. *J. Business* **64**(4), 549–571 (1991)
 69. Peiers, B.: A high-frequency study on the relationship between central bank intervention and price leadership in the foreign exchange market. Unpublished manuscript, Department of Economics, Arizona State University, 1994
 70. Petersen, M.A., Fialkowski, D.: Posted versus effective spreads, good prices or bad quotes? *J. Financial Econ.* **35**(3), 269–292 (1994)
 71. Pictet, O.V., Dacorogna, M.M., Müller, U.A., Olsen, R.B., Ward, J.R.: Real-time trading models for foreign exchange rates. *Neural Network World* **2**(6), 713–744 (1992)
 72. Poterba, J.M., Summers, L.H.: Mean reversion in stock prices: evidence and implications. *J. Financial Econ.* **22**, 27–59 (1988)
 73. Roy, A.D.: Safety first the holding of assets. *Econometrica* **20**, 431–449 (1952)
 74. Subrahmanyam, A.: Risk aversion, market liquidity, and price efficiency. *Rev. Financial Studies* **4**, 417–441 (1991)
 75. Surajaras, P., Sweeney, R.J.: *Profit-making speculation in foreign exchange markets*. Polit. Econ. Global Interdependence. Boulder: Westview Press 1992
 76. Suvanto, A.: Foreign exchange dealing. ETLA, the Research Institute of the Finnish Economy, pp. 1–132, 1993
 77. Svensson, L.E.: An interpretation of recent research on exchange rate target zones. *J. Econ. Perspectives* **6**(4), 119–144 (1992)
 78. Tauchen, G.E., Pitts, M.: The price variability-volume relationship on speculative markets. *Econometrica* **51**, 485–505 (1983)
 79. Taylor, M.P., Allen, H.: The use of technical analysis in the foreign exchange market. *J. Intern. Money Finance* **11**, 304–314 (1992)
 80. Taylor, S.J.: *Modelling financial time series*. Chichester: J. Wiley & Sons 1988
 81. Wasserfallen, W., Zimmermann, H.: The behavior of intra-daily exchange rates. *J. Banking Finance* **9**, 55–72 (1985)
 82. Zhou, B.: Forecasting foreign exchange rates subject to de-volatilization. Working Paper. MIT Sloan School **3510**, pp. 1–24, 1994