FRACTALS AND INTRINSIC TIME – A CHALLENGE TO ECONOMETRICIANS

U. A. MÜLLER, M. M. DACOROGNA, R. D. DAVÉ,
O. V. PICTET, R. B. OLSEN AND J. R. WARD *

UAM.1993-08-16

June 28, 1995

A preprint by the O&A Research Group

Presented in an opening lecture of the XXXIXth International Conference of the Applied Econometrics Association (AEA), *Real Time Econometrics — Submonthly Time Series*, 14-15 Oct 1993, in Luxembourg, and at the 4th International PASE Workshop, 22-26 Nov 1993, in Ascona (Switzerland)

Contents

* **Olsen & Associates**, Research Institute for Applied Economics, Seefeldstrasse 233, CH–8008 Zürich

The authors would like to thank the participants of the AEA conference in Luxembourg (14-15 Oct 1993) and the PASE Workshop in Ascona (22-26 Nov 1993) for their comments and help in correcting some errors.

Abstract

A fractal approach is used to analyze financial time series, applying different degrees of time resolution, and the results are interrelated. Some fractal properties of foreign exchange (FX) data are found. In particular, the mean size of the absolute values of price changes follows a "fractal" scaling law (a power law) as a function of the analysis time interval ranging from a few minutes up to a year. In an autocorrelation study of intra-day data, the absolute values of price changes are seen to behave like the fractional noise of Mandelbrot and Van Ness rather than those of a GARCH process.

Intra-day FX data exhibit strong seasonal and autoregressive heteroskedasticity. This can be modeled with the help of new time scales, one of which is termed intrinsic time. These time scales are successfully applied to a forecasting model with a "fractal" structure for FX as well as interbank interest rates, the latter presenting market structures similar to the Foreign Exchange.

The goal of this paper is to demonstrate how the analysis of high-frequency data and the finding of fractal properties lead to the hypothesis of a heterogeneous market where different market participants analyze past events and news with different time horizons. This hypothesis is further supported by the success of trading models with different dealing frequencies and risk profiles. Intrinsic time is proposed for modeling the frame of reference of each component of a heterogeneous market.

1 Introduction

For some years now, the set of available data from financial markets has grown rapidly. In the seventies, most of the empirical studies were based on yearly, quarterly, or monthly data. This data could typically be modeled by random-walk or linear models such as ARIMA. During the eighties, the study of weekly and daily data led to the discovery of new, non-linear properties, mainly autoregressive heteroskedasticity. Non-linear models such as ARCH, GARCH, and their variations were developed for this.

The first studies of intra-daily data in the nineties (Baillie and Bollerslev, 1990; Engle et al., 1990; Müller et al., 1990; Goodhart and Figliuoli, 1991; Dacorogna et al., 1993) revealed a new wealth of properties such as daily seasonal heteroskedasticity which cannot be fully covered by ARCH variations. Thanks to the enormous number of intra-daily observations, the confidence limits of statistical tests have become narrower, thus rejecting simple models that still succeed in explaining daily or weekly time series. The problem of aggregation arises: how can the properties of weekly price changes be covered by a model equation that focuses on a series of hourly price changes?

Mandelbrot introduced the *fractal* model to describe a certain class of objects exhibiting a complex behavior. He first applied it to financial data in (Mandelbrot, 1963). The fractal view as presented in (Mandelbrot, 1983) starts from a basic principle: analyzing an object on different scales, with different degrees of resolution, and comparing and interrelating the results. For time series, this means using different "time yardsticks", from hourly through daily to monthly and yearly, within the same study. This is far from the conventional time series analysis, which focuses on regularly spaced observations with a fixed time-interval size.

Some fractal properties can indeed be demonstrated by a study of high-density foreign exchange (FX) data. The mean size of the absolute values of price changes follows a "fractal" scaling law, a power of the observation time-interval size (Müller et al., 1990). In an auto-correlation study with high-density data (Dacorogna et al., 1993), the absolute values of price changes behave like the "fractional noise" of (Mandelbrot and Van Ness, 1968) rather than the absolute price changes expected from a GARCH process: the memory in the volatility declines hyperbolically with time. On the other hand, the nature of FX data is more complicated than a regular, self-similar fractal; this can be demonstrated by other empirical studies.

A precondition for analyzing intra-day data is the understanding and modeling of the strong daily *seasonal* heteroskedasticity. To do this, a new time scale ϑ is introduced. This time scale models the intra-daily and intra-weekly seasonality and is based on an analysis of the activity of the geographical centers of the worldwide, twenty-four hour FX market. Active periods are expanded and inactive ones, such as weekends, are contracted. Thus, ϑ-time is a kind of business time.

A second time scale τ is also proposed to model the non-seasonal fluctuations of the volatility. The idea of expanding volatile periods and contracting inactive ones is similar to concepts defined in (Mandelbrot and Taylor, 1967), (Allais, 1974), (Stock, 1988). This τ-time is called *intrinsic time* as it is determined by the time series itself rather than by an external clock. It is applied to a FX forecasting model with a "fractal" structure, where the model equations for short-term and long-term forecasts share the same structure on different scales. The same type of model also successfully works for interest rate forecasts.

The success of the fractal approach gives rise to the hypothesis that the market itself is "fractal", with *heterogeneous* trading behaviors. We contend that the different market participants have different time horizons in their analysis of past events and news and in their trading goals. To support this hypothesis further, we show that real-time trading models with different trading frequencies, developed and tested with a technology described in (Pictet et al., 1992), can be profitable at the same time. The market is composed of different actor groups or components, each with its own time horizon or frame of reference. It is driven by the actions of all these components and can thus be described as "fractal". The frame of reference of each component might be modeled in form of an intrinsic time of the component.

The empirical findings presented in this paper are explained with a minimum of detailed information. Many detailed questions, such as data filtering, are important in high-density data research but their presentation here would divert the attention from the main results. The discussion of details can be found in the literature.

In section 2 of this paper, some elementary fractal properties of FX rates are discussed. The ϑ-time scale is introduced in section 3 to model the daily and weekly seasonality and used for deseasonalization in section 4, where an autocorrelation study of volatility leads to the heterogeneous market hypothesis proposed in section 5. The definition of intrinsic time τ is given in section 6 and applied to the forecasting study presented in section 7. The results of this study and the analysis of trading models in section 8 further support the heterogeneous market hypothesis. Some implications of these results are discussed in the conclusions in section 9.

2 The viewpoint of fractal theory

The fractal approach to analyzing objects of different kinds, including time series, can be formulated as follows:

> **Objects are analyzed on different scales, with different degrees of resolution, and the results are compared and interrelated.**

In the introduction to the concept of fractals in (Mandelbrot, 1983), an example is presented that was first analyzed by Richardson. The length of the coastline of Great Britain is measured with yardsticks of different sizes. The length of the coastline is a function of the size of the yardstick: the smaller this size, the longer the total length measured because smaller yardsticks can follow more details of the coastline. The total length as a function of the size of the yardstick follows a power law.

If this principle is adopted to time series analysis, it means using different "time yardsticks", for example hourly, daily, weekly, or monthly ones. The scope of different possible analysis intervals is greatly increased when intra-day data is used. The real-time intra-day FX quotes published by market makers and distributed by information vendors such as Reuters, Knight-Ridder and Telerate are irregularly spaced in time, so there is no natural time grid for their analysis. A careful intra-day research project necessarily implies the use of more than one degree of time resolution.

The conventional time series analysis, focusing exclusively on a time series of regularly spaced observations, is far removed from both the fractal viewpoint and the real nature of the raw data: a regularly spaced economic time series is not original data but a preprocessed artifact. This is especially true for daily FX data: the choice of the time of day for recording daily observations from the twenty-four hour FX market is arbitrary.

A process equation that successfully explains daily price changes, for example, is unable to characterize the nature of hourly price changes. On the other hand, the statistical properties of monthly price changes are often not fully covered by a model for daily price changes. The new challenge to theoreticians is the development of consistent models that successfully and simultaneously characterize both the short and long-term behaviors of a time series.

A statistical study from the fractal viewpoint is based on analysis time intervals t of different sizes. An elementary example is the scaling law reported in (Müller et al., 1990), which relates the mean of the absolute logarithmic price change $|x|$ to the time-interval size t over which the price change is observed:

$$\overline{|x|} = \left(\frac{t}{T}\right)^D \tag{2.1}$$

where the bar over $|x|$ denotes the mean over a long sample, T is an empirical time constant and D the empirical drift exponent. The logarithmic price x is defined as follows:

$$x = \frac{\log(p_{\text{bid}}) + \log(p_{\text{ask}})}{2} \tag{2.2}$$

with the bid and ask prices, p_{bid} and p_{ask}.

This empirical scaling law has been studied with high-density FX data collected from raw data vendors such as Reuters, Knight-Ridder and Telerate. The result is presented in Figure 1 as a remarkably straight line on a double-logarithmic scale over a wide range of analysis time intervals. The ratio between the shortest analysis interval (10 minutes) and the longest one (1 year) in Figure 1 is greater than 50000. No structural break between the intra-day and the long-term behavior can be found.

The Wiener process, a continuous Gaussian random walk, also exhibits a scaling law like equation 2.1 with a drift exponent of 0.5. The empirical values of drift exponents D for

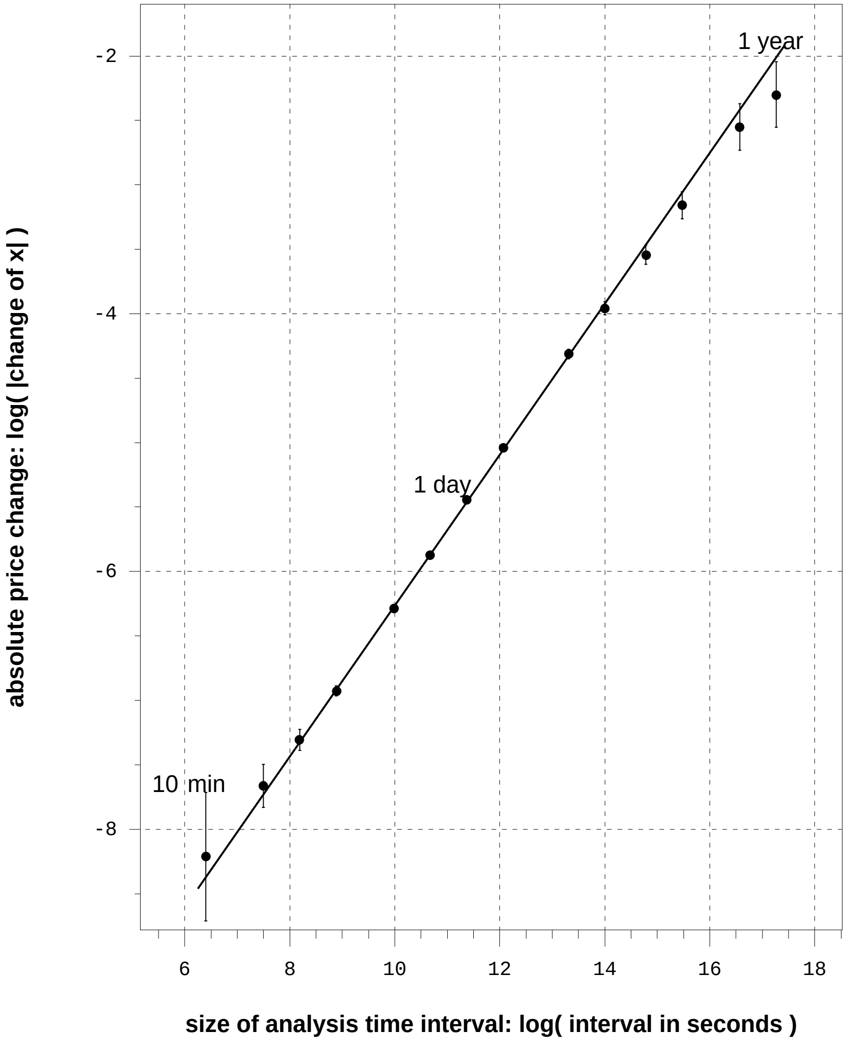

The mean absolute change of the logarithmic price x (defined by equation 2.2), plotted against the time-interval size over which it has been observed. FX rate: USD-DEM; sampling period: from Feb 1986 to Sep 1993. Double-logarithmic plot. The vertical bars indicate the standard errors of the observations. Resulting parameters of equation 2.1: drift exponent $D = 0.586$ with a standard error of 0.005, $T =$ 11230 days.

Figure 1: Scaling law of absolute price changes

different freely floating FX rates (as, for example, USD-DEM[1] in Figure 1) are clustered around a significantly higher value of 0.59. For the rates of the European Monetary System (EMS), the drift exponent is lower than 0.5. For the DEM-NLG rate, the most stable between two EMS currencies, a drift exponent of only 0.24 was measured. A striking example is the DEM-ITL rate. In the five years from June 1987 to May 1992, when the Italian Lira (ITL) was in the EMS, a drift exponent of 0.48 was found. For the year from Oct 1992 to Sep 1993, after the Lira had left the EMS, we measured a drift exponent of $D = 0.59$.

In spite of its elementary nature, a scaling law study is immediately able to reject the Gaussian random walk hypothesis and reveal an important property of financial time series: freely floating markets can be distinguished in their statistical behavior from regulated markets.

The most regular type of fractal is termed *self-similar* in (Mandelbrot, 1983). This means that the (statistical) properties of an object, analyzed with yardsticks of different sizes, are the same except for a scaling factor which is a power function of the size of the yardstick. If only the scaling law (Figure 1) is considered, a FX time series can be called a self-similar fractal.

[1]USD = US Dollar, DEM = German Mark. Throughout this paper, the ISO standard codes are used.

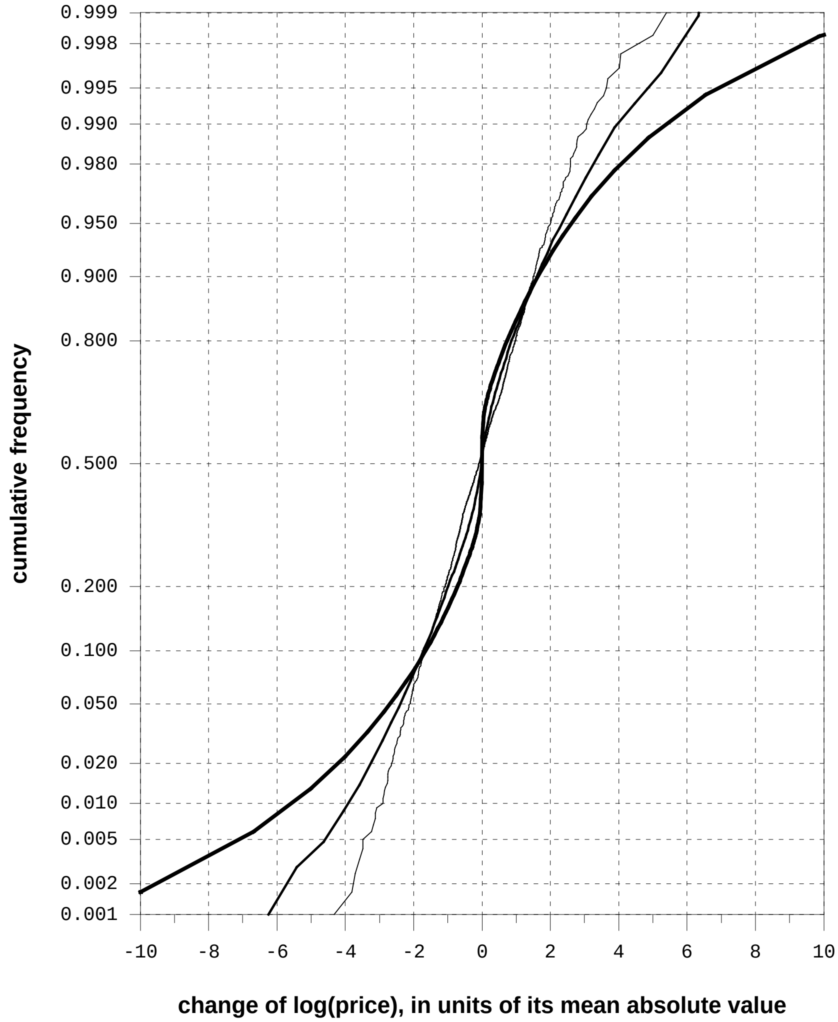

Cumulative distribution functions of price changes, observed over time intervals of 30 minutes (bold curve), 1 day (curve in the middle), and 1 week (thin curve). FX rate: USD-DEM; sampling period: from 5 May 1986 to 4 May 1993. The abscissa is in units of the mean absolute price change $\overline{|x|}$ measured in the same sampling period. For 30-minute intervals, $\overline{|x|} = 0.00045$, for daily intervals, $\overline{|x|} = 0.00414$, and for weekly intervals, $\overline{|x|} = 0.0124$. The cumulative frequency is plotted in such a way as to make any normal distribution appear as a straight line.

Figure 2: Distribution forms

There are, however, other statistical properties that can be compared between different analysis time intervals. In the case of self-similarity, the distribution function of price changes should have the same form for all interval sizes. The time series of cotton prices analyzed in (Mandelbrot, 1963) is self-similar, at least when daily and monthly price changes are compared: both have the same, stable distribution form. Figure 2 shows that FX rates do not share the same stability. Their distribution function becomes increasingly fat-tailed with higher time resolution (smaller yardsticks)[2]. Thus, FX rates are not self-similar fractals. Their fractal nature is more complicated and requires a deeper analysis.

There is yet another property of FX rates different from self-similar fractals: the daily and weekly seasonality of absolute price changes.

[2]The results are for USD-DEM, but all other major FX rates behave similarly.

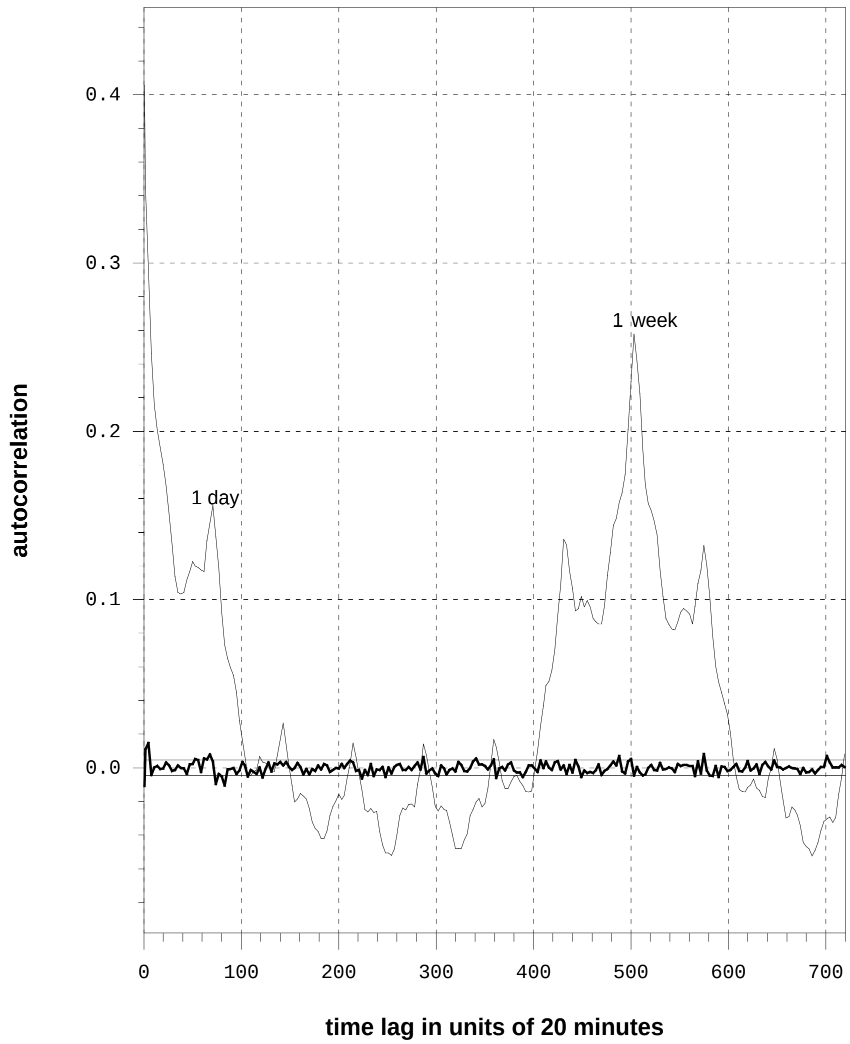

Autocorrelation of 20-minute price changes (bold curve) and the absolute values of the price changes (thin curve) computed on the physical time scale (uniform time, no special treatment of weekends and holidays). FX rate: USD-DEM; sampling period: from 2 June 1986 to 1 June 1993. The confidence limits represent the 95% confidence interval of a Gaussian random walk.

Figure 3: Autocorrelation of 20-minute price changes and their absolute values

3 A time scale to model the seasonality

The seasonality of a time series can be demonstrated by plotting an autocorrelation graph as in (Dacorogna et al., 1993) and in Figure 3. The bold curve in this figure indicates no significant autocorrelation of price changes x over intervals of 20 minutes. The corresponding autocorrelation function of the *absolute* values $|x|$ of the price changes, however, has a rich structure of significant peaks that indicate strong seasonality. The absolute value $|x|$ or its mean over a certain sample can be regarded as a basic measure of *volatility*. Therefore, the thin curve in Figure 3 indicates a strong seasonality of the volatility[3].

The autocorrelation in Figure 3 is analyzed with a basic time interval of only 20 minutes. The number of observations in seven years is thus very large and the significance limits, based on a Gaussian random process about the sample mean, very narrow. For lags of any integer number of days, clear peaks are found. They indicate the daily seasonality. The weekly seasonality is highly visible in the form of high autocorrelation for lags around one week and low, even

[3]The results are for USD-DEM, but all other major FX rates behave similarly.

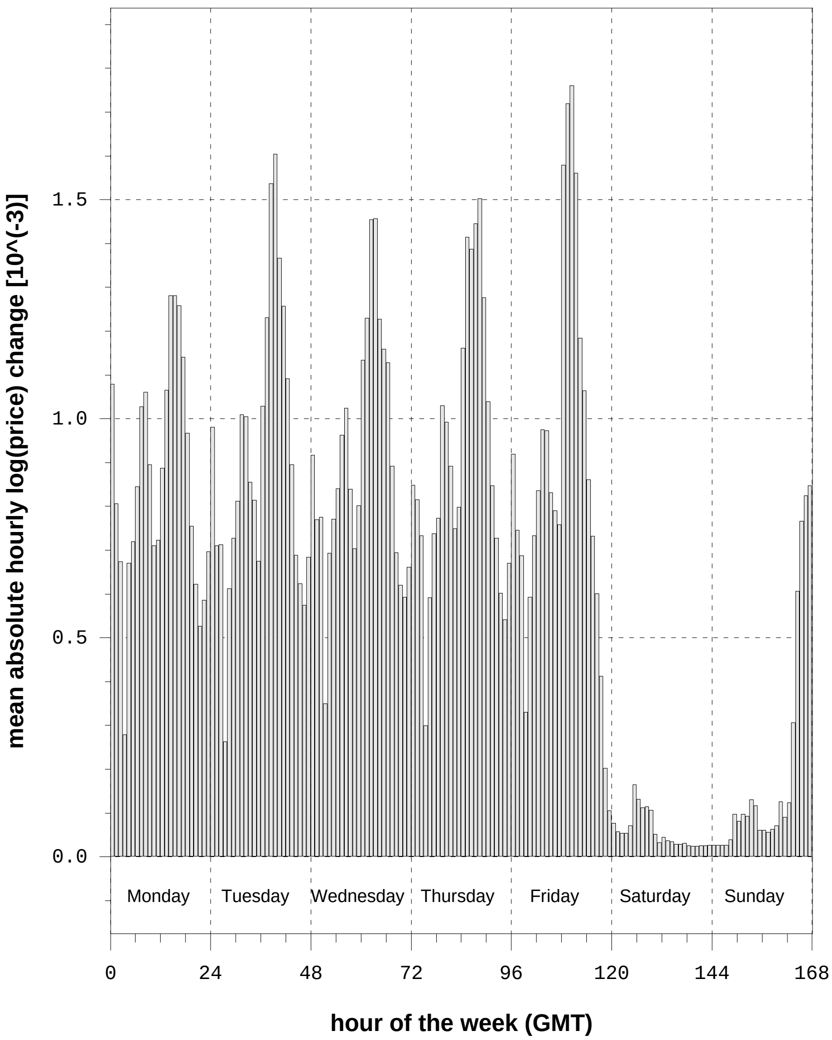

Average absolute hourly price changes. The hourly intra-week grid is in Greenwich Mean Time (GMT). No adjustment is made for daylight saving time. FX rate: USD-DEM; sampling period: from 3 Feb 1986 to 26 Sep 1993 (398 full weeks).

Figure 4: Intra-weekly histogram of mean absolute hourly price changes

negative autocorrelation for lags of about half a week (which frequently means the correlation of working days and weekends). Finally, there is a finer structure with small but visible peaks at integer multiples of eight hours, corresponding to a frequency three times the daily frequency. Apart from these seasonal peaks, there must be a positive component of the autocorrelation that declines with increasing lag. In Figure 3, this component cannot be analyzed yet as it is overshadowed by seasonality.

The seasonality can be demonstrated more directly by an intra-weekly analysis that allows us to study the daily and weekly seasonality. For this study, a sampling granularity of t = 1 hour is chosen. The week is subdivided into 168 hours from Monday 0:00 - 1:00 to Sunday 23:00 - 24:00 (Greenwich Mean Time, GMT) with index i. Each observation of the analyzed variable is made in one of these hourly intervals and assigned to the corresponding subsample with the correct index i. The 168 subsamples together constitute the full sample. The sample pattern is independent of bank holidays and daylight saving time. Any analyzed variable can be plotted as a histogram against the 168 hours of the statistical week, showing the typical intra-day and intra-week patterns of the variable.

An example of such an analysis is shown in Figure 4: the mean absolute hourly price change (a possible definition of volatility) in the statistical week[4]. The patterns of volatility are clearly correlated to the changing presence of the main market places of the worldwide FX market. The lowest volatility and market presence are observed on weekends. During the working week, the lowest volatility and market presence happen during the lunch hour in Japan (noon break in Japan, night in America and Europe). The highest values are reached in the early European afternoon, which coincides with the American morning. Many empirical studies give substantial evidence in favor of a positive correlation between price changes and volume in financial markets, as shown in (Karpoff, 1987). In (Baillie and Bollerslev, 1990), a similar intra-daily pattern of volatility as in Figure 4 is found with slightly different means.

The correlation of market presence and volatility leads to a central idea: modeling and explaining the empirically found seasonal volatility patterns with the help of fundamental information on the presence of the main markets around the world. We know the main market centers (e. g. New York, London, Tokyo), their time zones, and their usual business hours. When business hours of these market centers overlap, market activity must be attributed to their cumulative presence; it is impossible to assign the market activity to only one financial center at these times. The typical opening and closing times of different markets can be determined from our database, which also contains the originating locations of the quoted prices.

In (Dacorogna et al., 1993), the seasonal volatility patterns are modeled by a new *time scale* named ϑ-time. This time scale is modeled with the assumption of three main geographical areas where most of the worldwide trading activity is centered: East Asia, Europe, and America. This is in line with the small eight-hour seasonality found in Figure 5, corresponding to three market activity peaks in the 24 hours of a day. The ϑ-time scale expands times of day with a high mean volatility and contracts times of day with a low mean volatility and the weekends with their very low volatility. The weekends and business holidays are greatly reduced by the ϑ-scale. Most researchers of daily data already do the same thing: they omit the weekend days in their analysis. Hence, the ϑ-scale can also be termed a business time scale. The time scale expansions and contractions are defined in such a way as to eliminate the seasonality. The ϑ definition also makes use of the scaling law of equation 2.1, as explained in (Dacorogna et al., 1993). The ϑ-time is calibrated in such a way as to flow neither more slowly nor faster than physical time in the long-term average.

The strong daily and weekly seasonality observed in an analysis based on physical time virtually vanishes when analyzed in ϑ-time. The autocorrelation plot in Figure 5 demonstrates this[5]. The seasonal peaks have almost vanished; a steady, positive component of the autocorrelation becomes apparent. (At the lag of one week[6], a sharp, small remaining peak indicates a certain imperfection of the ϑ model). The deseasonalizing effect of ϑ-time can best be studied by comparing Figure 5 to Figure 3. A detailed discussion of Figure 5 follows in the next section.

[4] The results are for USD-DEM, but all other major FX rates behave similarly.

[5] The results are for USD-DEM, but all other major FX rates behave similarly.

[6] A week consists of 504 intervals of 20 minutes. A lag of a full business week in ϑ-time is slightly larger (522 intervals) in Figure 5 because the ϑ computation algorithm and the ϑ calibration take business holidays into account, see (Dacorogna et al., 1993).

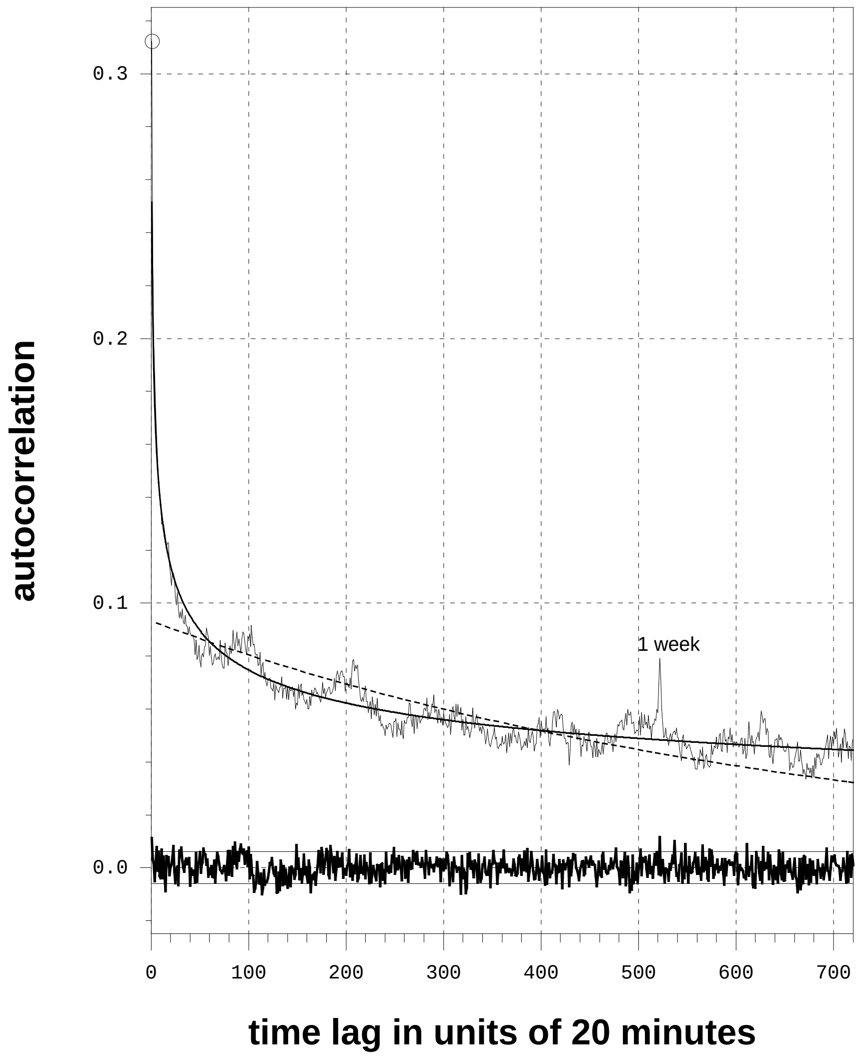

Autocorrelation of 20-minute price changes (bold curve) and the absolute values of the price changes (thin curve) computed on the ϑ-time scale. FX rate: USD-DEM; sampling period: from 5 May 1986 to 4 May 1992. A circle indicates the autocorrelation of absolute price changes at lag 1. A hyperbolic function (solid curve) and an exponential function (dotted curve) are shown, both representing the best fit of the autocorrelation of absolute price changes. The confidence limits represent the 95% confidence interval of a Gaussian random walk.

Figure 5: Autocorrelation of 20-minute price changes and their absolute values, ϑ-time

4 Volatility of price changes – a fractional noise?

The autocorrelation study shown in Figure 5 is based on FX data deseasonalized by using the ϑ-time, equally spaced by intervals of 20 minutes on the ϑ-scale. During active market periods, a ϑ-time interval of 20 minutes is shorter than a 20-minute interval in physical time; during inactive periods, it is longer. The choice of 20 minutes is arbitrary: the use of hourly steps, for example, yields similar results.

The autocorrelation of price changes x (bold curve in Figure 5) is still insignificant, as in Figure 3. The autocorrelation study of absolute price changes $|x|$ and hence volatility reveals two facts, as already shown in (Dacorogna et al., 1993):

1. In Figure 5, an exponential and a hyperbolic curve are plotted. Both curves represent the best fit of the autocorrelation curve. The comparison shows that assuming a hyperbolic decline is a good approximation; an exponential decline, as is typical for low-order ARCH and GARCH models, cannot be observed. The "memory" of the volatility is unexpectedly

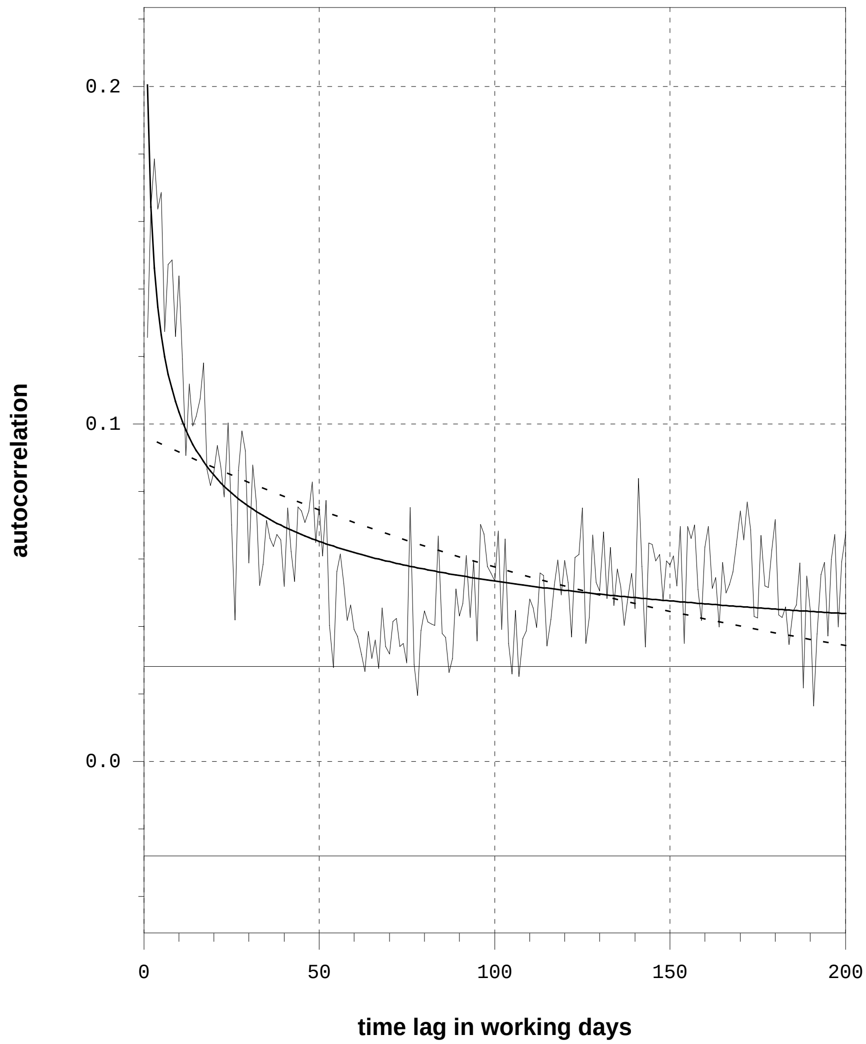

Autocorrelation of the absolute values of working-daily price changes analyzed with the help of the ϑ-time scale. FX rate: USD-DEM; sampling period: from 1 June 1973 to 1 June 1993. A hyperbolic function (solid curve) and an exponential function (dotted curve) are shown, both representing the best fit of the autocorrelation of absolute price changes. The confidence limits represent the 95% confidence interval of a Gaussian random walk.

Figure 6: Autocorrelation of absolute working-daily price changes

long: absolute 20-minute price changes are significantly correlated even if they occurred many days apart.

2. The residual deviation of the autocorrelation function from the pure hyperbolic fit exhibits a certain "heat wave" effect as defined by (Engle et al., 1990). This term was introduced in that paper for a continent-specific effect as opposed to "meteor shower" for an effect on the whole globe. At time lags of about one or two business days[7], indicating the presence of the same market participants, the residual autocorrelation is higher than at lags of one half or one and a half business days[8], when different market participants on opposite sides of the globe are active.

The autocorrelation of absolute price changes behaves like a fractal in a specific sense: in Figure 6, we find the same hyperbolic long-memory property for absolute working-daily price changes[9] as for 20-minute intervals in Figure 5. Although the analysis of Figure 6 extends over

[7]about 104 and 209 intervals of 20 minutes, see also the footnote at the end of section 3.
[8]about 52 and 157 intervals of 20 minutes, see also the footnote at the end of section 3.
[9]The results are for USD-DEM, but all other major FX rates behave similarly.

20 years of daily observations, there are still many more 20-minute observations in the six-year sample of Figure 5 than daily observations in the sample of Figure 6. The autocorrelation curve of Figure 6 is therefore noisier, with broader confidence limits.

A long memory of the volatility process was recently also found in (Ding et al., 1993) for daily stock index data. In that paper, the autocorrelation of different powers of absolute price changes, $|x|^\delta$, is found most significant when the exponent δ is approximately one. This finding supports our choice of absolute price changes (with $\delta = 1$) as the variable to study the autocorrelation of volatility in Figures 3, 5, and 6. The autocorrelation curves of (Ding et al., 1993) are also found to decline essentially more slowly than exponentially for large lags; a complicated function has been found to fit this decline. Without discussing the A-PARCH process developed in (Ding et al., 1993) (which cannot explain the form of the autocorrelation function empirically found in the same paper), we agree with the authors that a pure ARCH or GARCH process does not sufficiently describe the long memory of the volatility process.

There is a known theoretical process that has an asymptotically hyperbolic autocorrelation: the "fractional noise" of (Mandelbrot and Van Ness, 1968), which is a purely self-similar fractal. In (Mandelbrot, 1972), page 262, the autocorrelation function of fractional noise is given:

$$a \;=\; \frac{|l+1|^{2H} - 2\, l^{2H} + |l-1|^{2H}}{2} \tag{4.1}$$

where l is the autocorrelation lag and H the Hurst exponent, which lies between 0.5 and 1 for "persistent" fractional noise. For large lags l, the autocorrelation function converges to

$$a \;\approx\; H\,(2\,H-1)\; l^{2(H-1)} \tag{4.2}$$

which implies a hyperbolic decline. The autocorrelations of absolute price changes in Figures 5 and 6 also follow a hyperbolic decline. The introduction of the ϑ-time has thus helped to find a further fractal property of FX data. The exponent $2(H-1)$ of equation 4.2 can be empirically determined: we obtain the Hurst exponents $H = 0.87$ in Figure 5 and $H = 0.86$ in Figure 6. Inserting these H values, we can compute the factor $H(2H-1)$ of equation 4.2: we obtain 0.64 and 0.62 respectively. This factor is empirically found to be much lower: 0.25 and 0.20 respectively. This means that absolute price changes do not follow a *pure* fractional noise process. Absolute price changes are positive definite and have a skewed and fat-tailed distribution whereas the distribution function of pure fractional noise is Gaussian.

In (Peters, 1989) and (Peters, 1991), the existence of fractional noise in the price changes (rather than in their absolute values) is indirectly concluded from a R/S analysis that leads to a scaling law similar to equation 2.1. The author claims that a drift exponent different from 0.5 necessarily indicates fractional noise. This conclusion holds only if the distribution forms are stable, but Figure 2 shows that they are unstable, at least for submonthly analysis intervals. In our direct autocorrelation study of price changes (the bold curves in Figures 3 and 5), the fractional noise hypothesis is clearly rejected: unlike their absolute values, the price changes themselves exhibit no significant autocorrelation.

5 The hypothesis of a heterogeneous market

In Figure 5, we found two effects: a hyperbolic decline of the autocorrelation and the "heat wave" effect explained in section 4. They were found in the absolute price changes and hence the volatility rather than in the price changes themselves. This confirms the importance of volatility, which has been identified in many recent studies as a central variable for describing the state of a market. Volatility characterizes the market behavior more deeply than merely giving an indication of the size of current or recent price movements. It is the visible "footprint" of less observable variables such as market presence and also market volume (for which information is hardly available in FX markets).

The fact that, contrary to traditional beliefs, volatility is found to be positively correlated with market presence, activity, and volume in (Karpoff, 1987), (Baillie and Bollerslev, 1990), and (Müller et al., 1990) also emphasizes the key role of volatility for understanding market structures. The serial correlation studies by (LeBaron, 1992a) and (LeBaron, 1992b) show that subsequent price changes are correlated in low-volatility periods and slightly anti-correlated in high-volatility periods. In continuous samples mixed from both low-volatility and high-volatility periods, this important effect indicating the forecastability of price changes cannot be seen; it is conditional to volatility. Thus, volatility is also an indicator for the persistence of trends.

These recently found properties of volatility lead us to the hypothesis of a heterogeneous market as opposed to the assumption of a homogeneous market where all participants interpret news and react to news in the same way. The heterogeneous market hypothesis is characterized by the following interpretations of the empirical findings:

1. Different actors in the heterogeneous market have different time horizons and dealing frequencies. On the side of high dealing frequencies are the FX dealers and market makers (who usually have to close all their open positions before the evening); on the side of low dealing frequencies are the central banks, commercial organizations, and, for example, the pension fund investors with their currency hedging. The different dealing frequencies clearly mean different reactions to the same news in the same market. The market is heterogeneous, with a "fractal" structure of the participants' time horizons as it consists of short-term, medium-term and long-term components. Each such component has its own reaction time to news, related to its time horizon and characteristic dealing frequency. If we assume the memory of volatility of one component to be exponentially declining with a certain time constant (as in a GARCH(1,1) process), the memory of the whole market is composed of many such exponential declines with different time constants. The superposition of many exponential declines with widely differing time constants comes close to a hyperbolic decline. There is an analogy in physics: the secondary radioactivity of a compound material, where many different radioactive isotopes with different time constants decay simultaneously.

2. In a homogeneous market, the more agents are present, the faster the price should converge on the "real market value", on which all agents with a "rational expectation" agree. Thus, the volatility should be negatively correlated with market presence and activity. In a heterogeneous market, different actors are likely to settle for different prices and decide to execute their transactions in different market situations. In other words, they create volatility. This is reflected in the empirically found positive correlation of volatility and market presence.

3. The market is also heterogeneous in the geographic location of the participants. This

immediately explains the "heat wave" effect. In section 4, we found that the memory in the volatility process is relatively weak at time lags of about one half or one and a half business days, when market actors on opposite sides of the globe are related to each other, and relatively strong at time lags of about one or two business days, when identical groups of participants are considered.

The market participants of the heterogeneous market hypothesis also differ in other aspects beyond the time horizons and the geographic locations: they can have different degrees of risk aversion, institutional constraints, and transaction costs. Further evidence in favor of the heterogeneous market hypothesis is given in sections 7 and 8.

6 Intrinsic time: a time scale to model the volatility

The daily and weekly seasonal aspect of volatility has been modeled by the ϑ-time introduced in section 3. The volatility also exhibits non-seasonal, autoregressive clusters, as can be seen in Figures 5 and 6 and in the ARCH literature.

For modeling the volatility in all its aspects, the introduction of intrinsic time is proposed. The intrinsic time τ is defined as the cumulated sum of a market activity variable which is a statistical measure of very recent volatility. The τ value at the j'th time series observation is defined as

$$\tau_j \equiv \tau_{j-1} + k \, \frac{\vartheta_j - \vartheta_{j-1}}{\vartheta_r} \, v_r^{1/D} \, T \qquad (6.1)$$

The last two factors together are inverse scaling law, equation 2.1, applied to a variable v_r, which is the recent volatility (not annualized); ϑ_r is a range parameter (the ϑ-time-interval size of the price changes considered for computing the recent volatility v_r). In the implementation of (Dacorogna et al., 1992), this volatility is defined quite simply as an absolute price change:

$$v_r = |x(\vartheta_j) - x(\vartheta_j - \vartheta_r)| \qquad (6.2)$$

where $\vartheta_r = 1$ hour is chosen to reflect a short-term volatility. The factor k is calibrated in such a way that τ-time flows neither more slowly nor faster than physical time or ϑ-time in the long-term average. In (Dacorogna et al., 1992), k is a constant gained from a statistical analysis. The underlying time scale of the τ definition is always ϑ-time rather than physical time, so that the time series is analyzed in its deseasonalized form.

The forecasting results presented in this paper are based on a definition of recent volatility v_r more subtle and complicated than equation 6.2, with a variable rather than constant factor k that slowly adapts to changes in the time series behavior. These new definitions are described in (Müller, 1992).

The τ-scale expands volatile periods and contracts inactive ones, like ϑ-time, which does the same but only for the seasonal fluctuations of volatility. The term "intrinsic time" is justified as it reflects the behavior of the time series itself rather than that of an external clock. The idea of introducing an alternative time scale is not new. In (Mandelbrot and Taylor, 1967), a "transaction clock" was proposed which ticks once at every transaction of the worldwide market. (Allais, 1974) and (Stock, 1988) proposed other alternative time scales.

As an intrinsic time, τ relies on the availability of data from the time series. Since the time series values of the future are not available, τ-time is known only for the past and the present, whereas ϑ-time is based on calendar information which is also available for the future. Modeling the behavior of the τ and thus the volatility requires more than a mere definition of τ: it requires a forecasting model for τ. Introducing τ-time in a forecasting model is an alternative to ARCH-type models as it models the same effect, conditional heteroskedasticity, with different means.

In the fractal view of the market, the volatility should be studied in all time resolutions. The volatility of hourly price changes can be substantially different from that of, say, weekly price changes. Thus, we can define different τ-scales with different "time yardsticks" ϑ_r. In a heterogeneous market model, each market component can be modeled with its own τ-scale, reflecting its own perception of recent volatility.

Some studies such as (LeBaron, 1992a) and (LeBaron, 1992b) demonstrate that volatility characterizes the market behavior more deeply than merely giving an indication of the size of recent price movements, as discussed in section 5. Consequently, the use of the volatility-based τ-time in a price forecasting model can be beneficial beyond modeling and forecasting the absolute size of price changes.

7 A forecasting model with intrinsic time and fractal treatment of time horizons

In (Dacorogna et al., 1992), a univariate forecasting model for FX rates based on intrinsic time τ is presented and explained. The basic structure of this forecasting model is fractal in a sense: different forecasting intervals are treated with individual, independent forecasting models, but all of them with an identically structured forecast equation. The forecast of a price change is based on a linear combination of indicators which are the results of non-linear statistical operators; for example, non-linearly mapped moving averages over the recent past of the time series. These indicators conceptually come from simple trading systems used in practice by market participants, see for example (Dunis and Feeny, 1989). These trading systems yield buy and sell signals by evaluating an indicator function: the crossing of a certain threshold by the indicator on the positive side is regarded as a buy signal and on the negative side as a sell signal. An indicator is thus used as a predictor of a variable or its change. Short-term indicators reflect the (simplified) view of market actors with a short time horizon and long-term indicators that of market actors with a long time horizon.

Each price forecast for a certain forecast time is a linear combination of indicators of different time horizons; this is the way in which the heterogeneous market hypothesis enters the forecasting model. The time horizons of all these indicators are, however, not too different from the forecast time interval for which they are used: short-term forecasts are based on very short-term to medium-term indicators, long-term forecasts, on medium-term to very long-term indicators.

The indicator computations are based on τ-time, which expands periods of high volatility and contracts those of low volatility, thus better weighting the relative importance of events to the market. At the same time, the memory of the indicators becomes dynamic: short in high-volatility periods and long in low-volatility periods.

In (Dacorogna et al., 1992) and in the studies presented in this paper, all indicators are computed on the same τ-scale. A fully consistent fractal approach would go one step further and take a τ-time based on short-term volatility for short-term indicators and a τ-time based on long-term volatility for long-term ones.

A forecast is made in two steps:

1. a forecast of the intrinsic time interval τ from the current time point to the forecast time point; this is equivalent to a volatility forecast;
2. the price forecast: a forecast of the price change x from the current time point to the forecast time point.

Both forecasts are made with the same technology. The indicators of the τ forecast are computed on the ϑ-time scale, those of the price forecast on the τ-time scale.

The coefficients of the linearly combined indicators of the forecast equation are continuously re-optimized by a modified linear regression method on a *moving sample*. Unlike an ever-growing sample, a moving sample can adapt to long-term changes of the market behavior as it drops the observations of the very distant past. The time ranges of the moving samples also have a fractal structure: short-term forecasts are re-optimized with a rather short moving sample, long-term ones with a very long moving sample.

There are several methods of measuring the success of forecasting models, some of which are presented in (Dacorogna et al., 1992). Here, we use two measures: the direction quality and the signal correlation, both of which do not rely on any assumption about the underlying process. The direction quality is the percentage of forecasts in the right direction. As an approximation, we define the significance level as the 95% confidence level of the random walk:

$$\sigma_{\mathrm{dir},95\%} \approx \frac{1.96}{2\sqrt{n}} \tag{7.1}$$

where n is the number of observations in the forecast test. The factor 2 comes from the assumption of an equal probability of having positive or negative signals. A FX rate forecast is successful if its direction quality is significantly above 50%, that is, above 50% $+\sigma_{\mathrm{dir},95\%}$. The signal correlation is the correlation coefficient between the price change and its forecast. A successful model has a signal correlation significantly greater than zero. The 95% confidence level of the random walk is

$$\sigma_{\mathrm{sig},95\%} \approx \frac{1.96}{\sqrt{n}} \tag{7.2}$$

The forecast quality is evaluated on very large test samples. The total sample analyzed in (Dacorogna et al., 1992) consists of two main parts: the *in-sample* data used for the development and parameter selection of the model and the *out-of-sample* data, which was not used in the model development. In (Dacorogna et al., 1992), the results obtained from both sample types are reported separately. In sample as well as out of sample, the results are significantly better than random.

FX rate	forecast horizon	direction quality	signal correl.	sig.	FX rate	forecast horizon	direction quality	signal correl.	sig.
USD-DEM	2h	51.9	+2.2	+	DEM-JPY	2h	52.1	+4.1	+
	4h	51.9	+4.0	+		4h	51.2	+3.5	+
	8h	52.3	+3.2	–		8h	51.4	+2.7	–
USD-JPY	2h	52.5	+3.2	–	GBP-JPY	2h	53.2	+5.3	+
	4h	52.6	+4.8	+		4h	53.1	+6.3	+
	8h	51.8	+3.4	+		8h	52.8	+7.4	+
GBP-USD	2h	51.5	+1.4	–	GBP-DEM	2h	54.6	+7.7	+
	4h	51.6	+3.9	+		4h	53.5	+4.2	+
	8h	50.6	+3.4	–		8h	53.2	+4.6	+
USD-CHF	2h	51.8	+0.7	–	DEM-CHF	2h	54.6	+6.1	+
	4h	51.8	+1.5	–		4h	54.0	+5.1	+
	8h	51.6	+2.6	–		8h	53.5	+6.9	+
XAU-USD	2h	54.1	+3.0	+	JPY-CHF	2h	52.4	+4.2	+
	4h	53.1	+3.4	+		4h	51.5	+4.4	+
	8h	52.9	+2.6	–		8h	51.4	+5.8	+

Table 1: Out-of-sample forecasting results in percent for 4 USD rates, the gold price, and 5 cross rates for the period from 3 Sep 1990 up to 2 Sep 1993. All direction qualities are above 50% and all signal correlations above 0%, both significantly so in all cases marked by "+". In the cases marked by "–", at least one of the two quality tests is insignificant.

The quality measures of a forecasting model similar to that of (Dacorogna et al., 1992) for some FX rates are presented in Table 1. The direction quality significance limits according to equation 7.1 are 50.9%, 51.2%, and 51.7% for the 2-hour, 4-hour, and 8-hour forecasts; the signal correlation significance limits according to equation 7.2 are 1.7%, 2.4%, and 3.4%. All the results in Table 1 are better than expected from a random walk, the majority of them significantly so. They were computed out of sample and the model parameters are identical for all FX rates.

For further evidence of the quality of this model, we tested it on the deposit rates used in the transactions between banks and collected from the same information vendors such as Reuters and Telerate. The market for this type of interest rates is very similar to the FX market and often composed of the same agents[10]. The interest rate forecasting model uses exactly the same structure and is not optimized: the FX model of Table 1 was applied to interest rates without any modification, that is, no interest rate data was used to "fit" the model. It is, therefore, justified to qualify the full sample from 5 January 1987 to 2 August 1993 as an out-of-sample test.

[10]In fact, there are many links between these two markets since all swap transaction prices are computed through this type of interest rates and the swap contracts are playing an increasing role in the FX market, see (International Monetary Fund, 1993; Bank for International Settlements, 1993).

interest rate	forecast horizon	direction quality	signal correlation	interest rate	forecast horizon	direction quality	signal correlation
USD-3m	12h	60.4	+20.2	USD-6m	12h	58.4	+9.6
	24h	55.9	+11.2		24h	54.2	+9.8
	48h	55.1	+8.2		48h	54.7	+10.0
DEM-3m	12h	61.5	+21.2	DEM-6m	12h	61.7	+21.4
	24h	58.0	+17.6		24h	57.2	+15.8
	48h	53.0	+5.7		48h	54.5	+15.2
JPY-3m	12h	61.0	+20.6	JPY-6m	12h	61.5	+20.1
	24h	56.9	+14.6		24h	56.6	+16.9
	48h	56.5	+12.0		48h	58.3	+18.8
GBP-3m	12h	58.3	+11.1	GBP-6m	12h	56.6	+14.6
	24h	55.6	+8.1		24h	54.9	+7.2
	48h	53.8	+7.6		48h	53.3	+6.1
CHF-3m	12h	59.2	+14.5	CHF-6m	12h	58.6	+19.2
	24h	55.4	+15.4		24h	55.7	+13.8
	48h	56.2	+12.8		48h	56.5	+10.9

Table 2: Interest rate forecasting results in percent for the period from 5 Jan 1987 up to 2 Aug 1993. Different currencies, different maturities (3m = 3 months, 6m = 6 months), different forecast horizons (from 12 to 48 hours).

The quality measures of the interest rate forecasts are presented in Table 2. The direction quality significance limits according to equation 7.1 are 51.4%, 52.0%, and 52.8% for the 12-hour, 24-hour, and 48-hour forecasts; the signal correlation significance limits according to equation 7.2 are 2.8%, 4.0%, and 5.6%. This means that all the results in Table 2 indicate a significant success of the model. They are even better than those for the FX rates in Table 1.

The study by (Dacorogna et al., 1992) is not the only one reporting successful FX rate forecasting. (Bekaert and Hodrick, 1992), for example, find "a strong positive persistence of the excess returns" of their FX model. These findings contradict the hypothesis of market efficiency in its traditional sense; this is discussed at the end of the following section.

8 Trading models in heterogeneous markets

The autocorrelation study presented in section 4 led to the hypothesis of a heterogeneous, fractal structure of the FX market in section 5. This hypothesis is also supported by the success of the forecasting models presented in section 7. The key element of this hypothesis is the co-existence of market actors with essentially different reactions to news and different trading strategies, in particular, different time horizons and dealing frequencies. The traders may be different, but

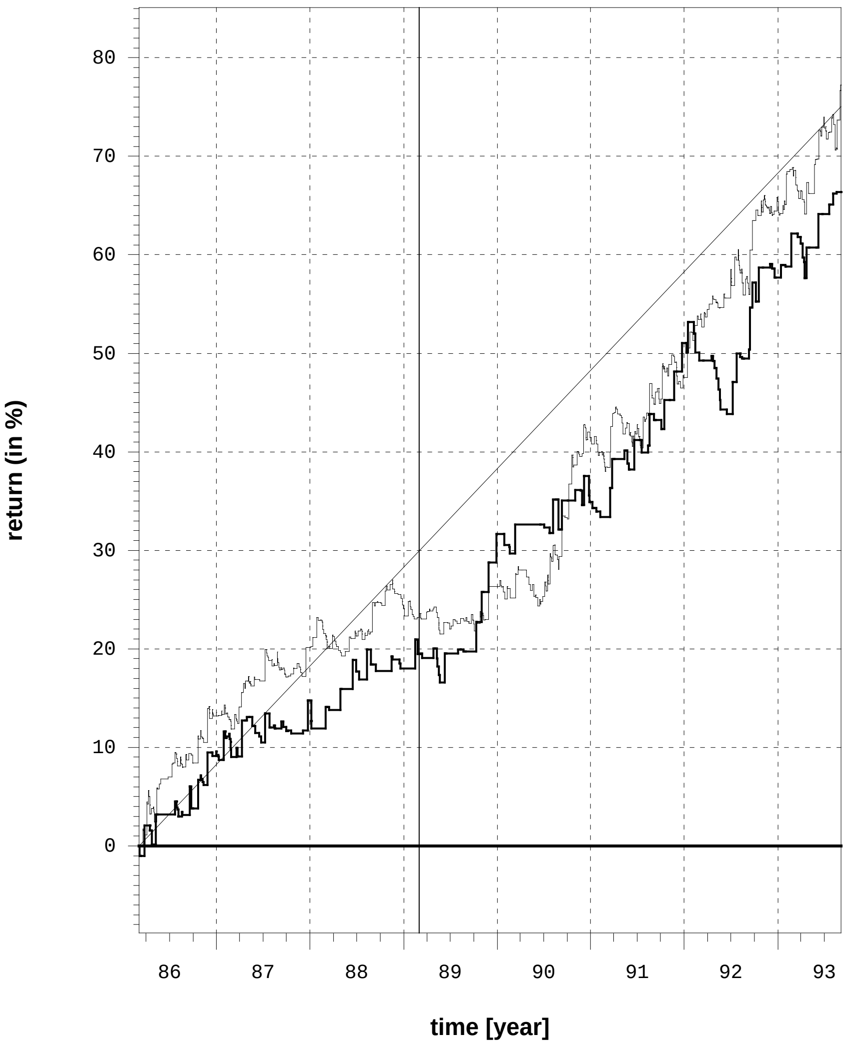

The returns of two trading models as functions of time. The two models are similar, but they have different time horizons and thus different average dealing frequencies: 2.9 transactions per month (bold curve) and 10.1 transactions per month (thin curve). FX rate: DEM-JPY. The results include transaction costs but no interest on any capital; profits are not re-invested. No leverage: the return figures are in percent of the maximum exposure limit. The straight line represents an annualized return of 10%; the vertical line separates the in-sample period used for optimization from the out-of-sample period used only for final testing.

Figure 7: Returns of two trading models with different dealing frequencies

they share the same goal of maximizing their profit expectation or, more precisely, their utility including a risk component (which is different for different traders in a heterogeneous market). The question arises whether traders can reach their goals even if their dealing frequencies are different.

It is possible to give further support to the heterogeneous market hypothesis by presenting two profitable real-time FX trading models based on intra-day data with very different dealing frequencies. Both these models use the technology described in (Pictet et al., 1992). The models give explicit trading recommendations under realistic constraints. They are only allowed to trade during the opening hours of a market, depending on the time zone and local holidays. The trading models presented in (Pictet et al., 1992) have proved successful not only in sample (with the data used for the model development), but also out of sample (where the data was used for final testing only) and, in particular, *ex ante* (the period after the last modification of the model).

The resulting behavior of the trading models is presented as a function of time in Figure 7. The return figures include the losses due to transaction costs (the bid-ask spreads). They assume an investor with a credit limit but no capital and, therefore, do not include any interest on any capital. The accumulated profits are not re-invested in our test.

The two univariate FX trading models presented in Figure 7 have been chosen as variations of the models presented in (Pictet et al., 1992) without any further optimization or parameter fitting. Nevertheless, both are profitable in the long run, with a moderate risk of temporary losses. The first model has a dealing frequency 3.5 times higher than that of the second model. It attempts to also profit from small, short-term price movements which the second model ignores. The dealing frequency of profitable trading can vary more than just by a factor of 3.5 as in Figure 7. The FX market makers can trade much more frequently than other market participants because the bid-ask spread is in their favor in each transaction.

The risk of temporary losses as in the loss periods of Figure 7 can be greatly reduced by diversification in a portfolio of models with different FX rates; see (Olsen & Associates, 1994).

If a certain profitable intra-day trading algorithm is tested with different business hours (corresponding to different time zones), its success can change considerably. Our systematic analysis has shown that the best choice of active trading hours is usually when the most important markets for the particular FX rate are active. All results demonstrate that the assumption of a homogeneous 24-hour FX market with identical dealers, following an identical "rational expectation", is far from reality.

Besides (Pictet et al., 1992), some other recent works have demonstrated the feasibility of profitable FX trading models, for example (Dunis and Feeny, 1989; Neftci, 1991). Both these papers deal with some phenomenological "technical analysis" methods employed by professional traders. In (Brock et al., 1992), the authors show that such technical trading rules can produce worthwhile returns even if transaction costs are included in the analysis. They conclude that past prices contain information on future price movements. A recent paper on the "active hedging" of foreign assets (Levich and Thomas, 1993a) shows a steady long-term profitability of a portfolio of FX trading models based on "simple technical trading rules". Other papers quoted in (Levich and Thomas, 1993a) also support this. In (Levich and Thomas, 1993b), the significant success of trading models based on technical trading rules is confirmed by extensive tests. In (Surajaras and Sweeney, 1992), the authors extensively discuss and analyze the possibility of consistent profit-making on the FX market with technical trading rules and conclude that on average the rules they tested could make about 3-4% per year on 15 major currencies over a test period of 2000 days. The authors of (Goodhart and Hesse, 1993) investigate some central bank interventions and show that these succeed in affecting FX rates in the right direction (in other words: gain some profits) if the return is measured a week or a month after the intervention. They also show that the success is unclear or even negative if measured only a day or less after the intervention. This finding is again in line with the heterogeneous market hypothesis: the central banks as a long-term component of the market optimize the long-term effect of their actions and give a lower weight to intra-day movements.

The profitability of trading models on the FX markets is not a particular case. Other financial markets show similar properties of consistent profit-making with trading models. The results on interest rate forecasting presented in section 7 are already an illustration of the fact that the market of interbank interest rates presents the same type of heterogeneity. We have recently developed successful trading models for US treasury bonds (Davé, 1994) based on a study of interest rate movements. The US treasury bond market involves other actors than the FX market but still presents many similarities, in particular, it is a very liquid market.

The stable profitability of trading models and the significant forecasting quality found in section 7 contradict the hypothesis of market efficiency in its traditional sense. A possible conclusion would be to identify the FX market as an inefficient market. A more constructive and appropriate alternative is to redefine market efficiency. The static efficiency definition that relies on instantaneous adjustment to news and a perfect, static market equilibrium at every moment might be replaced by a definition that looks at market dynamics; for example, at market frictions in form of transaction costs.

Even though some profitable trading strategies can be found, the FX market exhibits many properties related to efficiency: very high transaction volumes, a large number of market actors far from a monopoly situation, a more or less equitable access to information, a 24-hour market without business hour limitations, and, as a consequence of all these properties, low transaction costs (bid-ask spreads). A new definition of market efficiency might be based on the most directly quantifiable property: the size of the transaction costs. The simultaneous market presence of different types of traders with different time horizons and strategies may also turn out to be an important criterion for efficient markets, although in a new sense: it provides a wider and more diverse set of possible transaction partners to each individual trader.

9 Conclusion and outlook

In several studies of high-frequency intra-day FX data, different fractal properties of FX data have been observed: the scaling law and a behavior of absolute price changes which is similar to fractional noise. A fractal approach to price forecasting and the use of modified time scales that reflect market activity, such as ϑ-time and the intrinsic time τ, have been shown to be successful. Trading strategies with different dealing frequencies and other differences, for example in opening hours or risk aversion levels, can be successful at the same time.

As a common cause for these phenomena we propose the hypothesis of a heterogeneous market which assumes the existence of essentially different types of traders with a fractal structure of their time horizons, from short to long-term. The fact that market participants work on different continents with different time zones is also essential for the market dynamics, as demonstrated by a "heat wave" effect of the autocorrelation of absolute price changes (see Figure 5 and section 4). Further differences among market actors originate from institutional constraints. Most intra-day FX dealers, for example, can have open positions in foreign currencies only during the business hours of their time zone and have to close them in the evening.

The presence of various types of market participants is not limited to the FX market. We have shown that the market for interest rates lends itself to the use of similar approaches. More statistical studies on interest rates remain to be done to explore the existence of a scaling law or the behavior of the price change distributions but we are confident that, even if some statistical properties are not the same as for the FX rates, the "fractal" approach will provide useful insights into these markets, as already demonstrated by these first results.

A way to model the different types of market participants is to group them in market components. Each component has a view which is biased by its own trading characteristics and essentially different from all the others. It interprets and weights the same past of the time series and the same news differently, in its own frame of reference. As a consequence, we propose to model each market component with its own intrinsic time. This suggestion goes further than the use of τ-time presented in sections 6 and 7.

The heterogeneous market hypothesis as well as the systematic profitability of certain trading algorithms shown in section 8 violate the traditional "efficient market" hypothesis and thus raise some fundamental questions. Rather than rejecting market efficiency altogether, some ideas for a less static definition of efficiency are proposed in section 8.

The interaction of market actors following different strategies can also be simulated in the sense of experimental economics. A step in this direction is done in (Rust et al., 1993), where the interaction of different computerized traders in a simulated double auction market is studied. For future research, we suggest modeling the different traders of a market simulation study according to the heterogeneous market hypothesis: each market component with its own time horizon, geographic location, and institutional constraints. The forecasting study presented in section 7 can be seen as an early form of such a simulation project as it studies forecasts composed of indicators with different time horizons, representing different market components.

References

Allais M., 1974, *The psychological rate of interest*, Journal of Money, Credit and Banking, **3**, 285–331.

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

Bank for International Settlements ., 1993, *Central bank survey of foreign exchange market activity*, document from the Monetary and Economic Department, **Basle February 1993**, 1–42.

Bekaert G. and Hodrick R. J., 1992, *Characterizing predictable components in excess returns on equity and foreign exchange markets*, Journal of Finance, **47**(2), 467–509.

Brock W., Lakonishok J., and LeBaron B., 1992, *Simple technical trading rules and the stochastic properties of stock returns*, The Journal of Finance, **47**(5), 1731–1764.

Dacorogna M. M., Gauvreau C. L., Müller U. A., Olsen R. B., and Pictet O. V., 1992, *Short term forecasting models of foreign exchange rates*, Presentation at the IBM Summer Research Institute in Oberlech Austria on "Advanced Applications in Finance, Investment and Banking" held from July 27 to July 31, 1992. MMD.1992-05-12, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

Dacorogna M. M., Müller U. A., Nagler R. J., Olsen R. B., and Pictet O. V., 1993, *A geographical model for the daily and weekly seasonal volatility in the FX market*, Journal of International Money and Finance, **12**(4), 413–438.

Davé R. D., 1994, *Performance of US treasury security trading models*, Internal document RDD.1994-04-13, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

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

Dunis C. and Feeny M., 1989, *Exchange Rate Forecasting*, Woodhead-Faulkner, Cambridge.

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

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

Goodhart C. A. E. and Hesse T., 1993, *Central Bank Forex intervention assessed in continuous time*, Journal of International Money and Finance, **12**(4), 368–389.

International Monetary Fund ., 1993, *International capital markets. Part I. exchange rate management and international capital flows*, World Economic and Financial Surveys, **Washington, April 1993**, 1–79.

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

LeBaron B., 1992a, *Forecast improvements using a volatility index*, Journal of Applied Econometrics, **7**, 137–149.

LeBaron B., 1992b, *Some relations between volatility and serial correlations in stock market returns*, Journal of Business, **65**(2), 199–219.

Levich R. M. and Thomas L. R., 1993a, *Internationally diversified bond portfolios: the merit of active currency risk management*, working paper of the National Bureau of Economic Research, **4340**, 1–35.

Levich R. M. and Thomas L. R. III., 1993b, *The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach*, Journal of International Money and Finance, **12**(5), 451–474.

Mandelbrot B. B., 1963, *The variation of certain speculative prices*, Journal of Business, **36**, 394–419.

Mandelbrot B. B., 1972, *Statistical methodology for nonperiodic cycles: from the covariance to R/S analysis*, Annals of Economic and Social Measurement, **1**(3), 259–290.

Mandelbrot B. B., 1983, *The Fractal Geometry of Nature*, W.H.Freeman and Company, New York.

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

Mandelbrot B. B. and Van Ness J. W., 1968, *Fractional brownian motions, fractional noises and applications*, SIAM Review, **10**, 422–437.

Müller U. A., 1992, *Design of a new intrinsic time τ and its use in the Unbiased Forecaster (UBF)*, Internal document UAM.1992-05-18, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

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

Neftci S. N., 1991, *Naive trading rules in financial markets and Wiener-Kolmogorov prediction theory: a study of "technical analysis"*, Journal of Business, **64**(4), 549–571.

Olsen & Associates, 1994, *O & A Trading Model Portfolios, a Performance Analysis*, Release Documents. Olsen & Associates, Zurich.

Peters E. E., 1989, *Fractal structure in the capital markets*, Financial Analysts Journal, **1989**(July/August), 32–37.

Peters E. E., 1991, *Chaos and order in capital markets*, A Wiley Finance Edition. John Wiley & Sons, New York.

Pictet O. V., Dacorogna M. M., Müller U. A., Olsen R. B., and Ward J. R., 1992, *Real-time trading models for foreign exchange rates*, Neural Network World, **2**(6), 713–744.

Rust J., Miller J. H., and Palmer R., 1993, *Behavior of trading automata in a computerized double auction market*, in The Double Auction Market – Institutions, Theories, and Evidence; Santa Fe Institute Proceedings, **14**, 155–198.

Stock J. H., 1988, *Estimating continuous-time processes subject to time deformation*, Journal of the American Statistical Association, **83**(401), 77–85.

Surajaras P. and Sweeney R. J., 1992, *Profit-Making Speculation in Foreign Exchange Markets*, The Political Economy of Global Interdependence. Westview Press, Boulder.