

**TABLE 5.13** Correlation coefficients for activity measures.

Correlation coefficients computed for the different intraday analyses for the USD against DEM, JPY, CHF, and GBP and XAU (gold) against USD. Sampling period: from January 1, 1987, to December 31, 1993.

	DEM	JPY	GBP	CHF	XAU
$E( r )$ -ticks	+0.540	+0.421	+0.779	+0.755	+0.885
$E( r )$ -spread	-0.220	-0.485	-0.570	-0.704	-0.287
Ticks-spread	-0.693	-0.018	-0.881	-0.707	-0.450

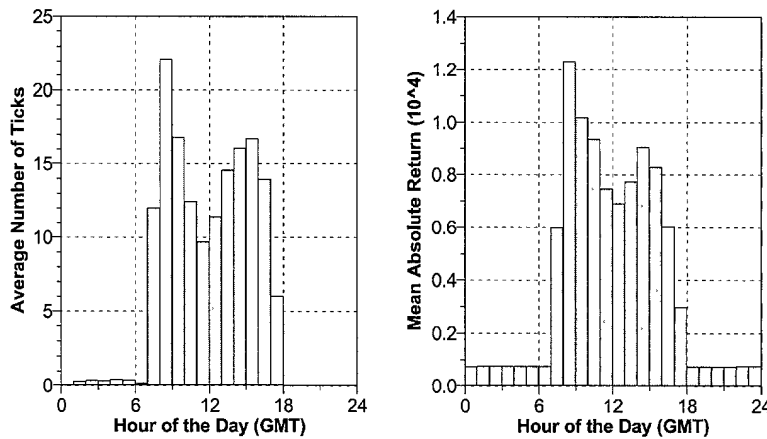
coefficients between the intradaily histograms of Figure 5.12 are also positive, as explicitly shown in the first line of Table 5.13. We conjecture that both variables are positively correlated to a third one, the worldwide intraday transaction volume, which is not known for the FX market. Transaction volume figures are, however, available for the stock market; their positive correlation to squared returns (and hence the volatility) has been found by Harris (1987) and other authors. Recently, Hasbrouck (1999) examined the data of the New York Stock Exchange and found similar correlations as in Table 5.13 for his transaction data, but the correlations did not uniformly increase when the data were aggregated.

The statistics show that an analysis of return distributions that neglects the large differences between the hours of a day and the days of the week is inappropriate. In Chapter 6, we will introduce a new time scale to solve this problem.

### 5.6.3 Seasonal Volatility: U-Shaped for Exchange Traded Instruments

Intradaily seasonalities were also found in the stock markets by Ghysels and Jasiak (1995), Andersen and Bollerslev (1997b) and Hasbrouck (1999). Unlike the FX market, stock exchanges and money market exchanges are active less than 24 hr a day. Thus the shape of the seasonality is different. It is called the U-shape because the high volatility of the opening is followed by a decrease, which is in turn followed by an increase of volatility just before closing. Ballochi *et al.* (1999b) study the Eurofutures markets and find the expected intraday seasonality. For all contracts traded on LIFFE the hourly tick activity displays the U-shape with its minimum around 11 a.m. to 1 p.m. (GMT) and a clustering of activity around the beginning and the end of the trading day. There are differences among Eurofutures between the levels and widths of the peaks and the level of the minimum. The Eurodollar (a contract type traded on CME, see section 2.4.1) displays similar behavior but the activity in the first half of the working day, which takes place when the European markets are still open, is higher than during the second half of the day, when European markets have already closed and Asian markets are not yet open.

REPRODUCED FROM




**FIGURE 5.13** Intraday analysis of Short Sterling in position two. The intraday tick activity (left histogram) displays the average number of ticks occurring in each hour of the day whereas the intraday volatility (right histogram) shows the mean absolute return. Both plots display similar U-shapes, the only difference being that the minimum appears one hour later for intraday returns. The time scale is GMT (not UKT, the local time used by LIFFE in London). The sampling period starts on January 1, 1994, and ends on April 15, 1997. The total number of ticks is 184,360.

Intraday returns follow a pattern similar to that presented by intraday tick activity. In general, opening hours show the highest price variation (the difference with respect to the average of the other hours is around one basis point); only in some cases does the largest return occur toward closing time (usually in the last positions). Differences occur in some positions<sup>28</sup> for Short Sterling, Eurolira, and Three-Month Ecu,<sup>29</sup> which display the minimum of the U-curve 1 hour later than in the tick-activity case. This can be seen in Figure 5.13, which displays intraday tick activity and intraday returns for Short Sterling in position two. Note that the U-shapes in this figure are blurred by the fact that Greenwich Mean Time (GMT) is used. The observations do not only cover winter months but also summers where the time scale used by LIFFE in London is shifted by 1 hour (daylight saving time). If the time scale was local time (UKT) instead of GMT, the U-shapes would be more pronounced with clearer peaks at opening and closing.

The first two positions of the Euromark display less regularity in the intraday return behavior. This behavior is confirmed also by correlation results: on the whole, the correlation between hourly tick activity and hourly returns is above 0.96; only Euromark for the first two positions and Three Month-Ecu for the fourth position show a lower correlation around 0.90. In general, for Eurodollar, Euromark and Short Sterling, hourly returns tend to increase from position 1

<sup>28</sup> For an explanation of the word “position,” see Section 2.1.2.

<sup>29</sup> Short Sterling, Eurolira, Three Month-Ecu, and Euromark are names of LIFFE contracts, all with an underlying 3-month deposit. Ecu is the European Currency Unit that preceded the Euro.



**FIGURE 5.14** Volatility as a function of time to expiry. The volatility values are daily averages over 36 contracts (9 for Eurodollar, 9 for Euromark, 9 for Short Sterling, and 9 for Eurolira). The abscissa corresponds to the time to expiry: the farther on the right-hand side, the farther away from expiry.

to position 4; Eurodollar and Short Sterling display a decrease for some hours in position 4.

Looking at intraweek tick activity, there is evidence of a day-of-the-week effect. In general, the level of activity displays a minimum on Monday and a maximum on the last two working days of the week, usually on Thursday for LIFFE contracts and on Friday for CME contracts. The difference is definitely significant for the Eurodollar; in fact, for positions 1 and 2 the tick activity on Friday is almost double that on Monday and it becomes more than double for positions 3 and 4. In general, there is a gradual increase from Monday to Friday.

#### 5.6.4 Deterministic Volatility in Eurofutures Contracts

Balocchi *et al.* (2001) provide evidence that the volatility of futures prices systematically depends on the time interval left until contract expiry. We call these systematic volatility patterns *deterministic*, as opposed to the also existing stochastic fluctuations of volatility. In order to probe the existence of a seasonality related to contract expiry, a sample consisting of many futures contracts is needed. For several Eurofutures contract type (Eurodollar, Euromark, Short Sterling, and Eurolira) and for each contract expiry, we build a series of hourly returns using linear interpolation. Then we compute daily volatilities taking the mean absolute value of hourly returns from 00:00 to 24:00 (GMT) of each working day (weekends and holidays are excluded). These daily volatilities are plotted against time to expiry. The result is shown in Figure 5.14. The vertical axis represents the mean volatility

computed from all Eurofutures and all contracts together. The horizontal axis represents the time left to expiry, as we move towards the left the number of days to expiry decreases.

Figure 5.14 spans a period of about 360 days because only within that period we are able to compute our mean volatility based on a full set of contracts. Some contracts have bad data coverage for times to expiry exceeding 360 days. The results obtained are quite interesting. There is a downward trend in volatility as the time to expiry decreases (moving from right to left in Figure 5.14). This downward trend is weak between about 300 and 180 days before expiry but becomes strong as we move toward the expiry date. There is also an unexpected behavior consisting of oscillatory movements with peaks every 90 days corresponding to rollover activities near the ending of contracts. These results are confirmed also by a deterministic volatility study on each single Eurofutures type—except Euroaira, which displays an increment in volatility as we move toward expiry. Eurodollar, Euromark, and Short Sterling show a decreasing volatility at least for the last 300 days before expiry. All Eurofutures display oscillatory movements with peaks around expiry dates (this appears particularly evident for Short Sterling).

A possible explanation for this effect is that these markets are all “cash settled” and therefore have no “delivery risk”; this means there is no risk of holding these futures on expiry day. Due to transaction costs, it is cheaper to take the cash at expiry than to close the position and realize the cash the day before. In other future markets such as the Deutsche Termin-Börse or the commodity markets, people who hold long positions to expiry actually take physical delivery of the underlying commodity or bond. There is a risk as expiry approaches as to which bond or type of commodity will be delivered. This may cause an increase in volatility as expiry approaches—a behavior opposite to that of Figure 5.14.

### 5.6.5 Bid-Ask Spreads

The bid-ask spread reflects many factors such as transaction costs, the market maker’s profit, and the compensation against risk for the market maker, see (Glassman, 1987; Glosten, 1987). The subject of the intraday and intraweek analysis is the relative spread  $s_j$  (Equation 3.12). It is usually below or around 0.1%, and its distribution is not symmetric. Negative changes are bounded as spreads are always positive, but the spread can exceed 0.5% in times of low market activity. The arithmetic mean of  $s_j$  weights these low-activity spreads too strongly and therefore we choose the *geometric* mean as a more appropriate measure:

$$\bar{s}_i = \left( \prod_{j=1}^{n_i} s_{i,j} \right)^{\frac{1}{n_i}} \quad (5.39)$$

The index  $i$  indicates the hour of a day (or a week) or the day of the week, depending on the analysis. The total number of ticks that belong to the  $i^{th}$  interval is  $n_i$ .  $j$  is the index and  $s_{i,j}$  the spread of these ticks.



**TABLE 5.14** Average spreads.

Geometric average of the relative spread for each day of the week (including weekends) for the USD against DEM, JPY, CHF, and GBP and XAU (gold) against USD; for the period from January 1, 1987, to December 31, 1993. The relative spread figures have to be multiplied by  $10^{-4}$ .

	DEM	JPY	GBP	CHF	XAU
Monday	4.57	5.72	4.82	6.32	12.58
Tuesday	4.52	5.64	4.77	6.28	12.51
Wednesday	4.57	5.71	4.81	6.32	12.49
Thursday	4.64	5.77	4.84	6.38	12.62
Friday	4.79	6.00	4.99	6.49	12.59
Saturday	7.69	17.91	17.32	18.02	13.26
Sunday	5.28	6.78	9.60	10.99	14.04


Müller and Sgier (1992) analyze in detail the statistical behavior of the quoted spread. Here we shall present their main conclusions. First, it is important to remember that all the statistical analyses are dominated by one property of quoted FX spreads, which is the discontinuity of quoted values (see Section 5.2.2). This data set contains price quotes rather than traded prices. The banks that issue these price quotes are facing the following constraints:

- Granularity: FX prices are usually quoted with five digits—that is, 1.6755 (USD-DEM) or 105.21 (USD-JPY). The lowest digit sets the granularity and thus the unit *basis points*.
- Quoted spreads are wider than traded spreads as they include “safety margins” on both sides of the real spread negotiated in simultaneous real transactions. These margins allow the FX dealers, when called by a customer during the lifetime of the quote, to make a fine adjustment of the bid and ask prices within the range given by the wide quoted spread. They can thus react to the most recent market developments.
- FX dealers often have biased intentions: while one of the prices, bid or ask, is carefully chosen to attract a deal in the desired direction, the other price is made unattractive by increasing the spread.
- Because quoted spreads are wider than traded spreads, they do not need the high precision required in the direct negotiation with the customer on the phone. Hence, there is a tendency to publish formal, “even” values of quoted spreads as discussed in Section 5.2.2.

The strong preference for a few formal spread values, mainly 5 and 10 basis points, clearly affects every statistical analysis.

The results are shown in the middle histograms of Figure 5.12 and in Table 5.14. The general behavior of spreads is opposite to those of volatility and tick frequency. Spreads are high when activity is low, as already noticed by

2025 RELEASE UNDER E.O. 14176



**FIGURE 5.15** Cumulative distributions of relative spreads (left) and logarithm of the relative spread (right) shown against the Gaussian probability on the y-axis. The distribution is computed from a time series of linearly interpolated spread sampled every 10 min for USD-DEM. The sample runs from March 1, 1986, to March 1, 1991.

Glassman (1987). FX spreads on Saturdays and Sundays can have double and more the size of those on weekdays and, as in Table 5.12, Sundays differ slightly less from working days than Saturdays. Sunday in GMT also covers the early morning of Monday in East Asian time zones. Unlike the volatilities, the average FX spreads exhibit a clear weekend effect in the sense that the Friday figures are higher, though still much lower than those of Saturday and Sunday. The spreads of gold vary less strongly, but they have double the size of the FX spreads on working days. The FX rate with the smallest spreads, USD-DEM, was the most traded one according to all the BIS studies (until it was replaced by EUR-USD in 1999). The histograms in Figure 5.12 have intraday patterns that are less distinct than those of volatility, but still characteristic. We analyze their correlations with both the volatilities and the numbers of quoted ticks. All the correlation coefficients on the second line of Table 5.13 and most of them on the third line are negative, as one would expect. The FX rates have different spread patterns. For USD-CHF, for instance, there is a general spread increase during the European afternoon when the center of market activity shifts from Europe to America, while the USD-JPY spreads decrease on average at the same daytime. This indicates that American traders are less interested in Swiss Francs and more in Japanese Yens than other traders. Hartmann (1998) uses the spreads to study the role of the German Mark and the Japanese Yen as “vehicle currencies,” as compared to the USD.

An analysis of the empirical cumulative distribution function of the relative spread  $s$  is shown in the left graph of Figure 5.15 for USD-DEM and for  $\ln s$  in the

right graph of Figure 5.15. The resulting cumulative distribution functions have the following properties:

1. They are not Gaussian, but convex ( $s$  strongly,  $\ln s$  slightly), indicating a positive skewness and leptokurticity (of the tail on the positive side).
2. They look like a staircase with smooth corners. For the nominal spread in basis point,  $s_{\text{nom}}$ , we would expect a staircase with sharp corners, the vertical parts of the staircase function indicating the preferred “even” values such as 10 basis points. Although  $s$  is a relative spread ( $\approx s_{\text{nom}}/p_{\text{bid}}$ ), where the bid price  $p_{\text{bid}}$  fluctuates over the 5-year sample, and although we use linear interpolation in the time series construction (see Section 3.2.1), the preferred “even”  $s_{\text{nom}}$  values are still visible.

0272

# 6

---

## MODELING SEASONAL VOLATILITY

### 6.1 INTRODUCTION

The intradaily and intraweekly seasonality of volatility is a dominant effect that overshadows many further stylized facts of high-frequency data. In order to continue the research for stylized facts, we need a powerful treatment of this seasonality.

Many researchers who study daily time series implicitly use, as a solution, a business time scale that differs from the physical scale in its omission of Saturdays, Sundays, and holidays. With the  $\vartheta$ -scale we extend this concept to the intraday domain, thereby allowing us to tackle a fundamental source of seasonality originating from the cyclical nature of the 24-hr hour trading around the globe in different geographical locations.

There are, therefore, three main motivations for our model:

- To provide a tool for the analysis of market prices by extending the concept of business time scale to intraday prices
- To make a first step toward formulating a model of market prices that also covers the intraday movements
- To gain insight into the interactions of the main market centers around the world and their relevance to each particular foreign exchange (FX) rate

A number of papers such as Andersen and Bollerslev (1997b, 1998b), Taylor and Xu (1997), and Beltratti and Morana (1999) propose alternative approaches for dealing with volatility seasonalities. They are based on a factorization of the volatility into an essentially deterministic seasonal part and a stochastic part, which is (more or less) free of seasonalities. The former is then modeled by a set of smooth functions. Cutting out the inactive periods of the time series and gluing together the active parts, Andersen and Bollerslev (1997b) succeeded in applying their method also to the S&P 500 index. This procedure is not fully satisfactory for a number of reasons: time series have to be preprocessed, there is no treatment of public holidays and other special days, the model fails when the opening or closing time of the market changes, and it is not adequate for instruments with a complex, hybrid volatility pattern. Gençay *et al.* (2001a) use the wavelet multiresolution methods for dealing with volatility seasonalities which is studied in Section 6.4.

## 6.2 A MODEL OF MARKET ACTIVITY

### 6.2.1 Seasonal Patterns of the Volatility and Presence of Markets

The behavior of a time series is called *seasonal* if it exhibits a periodic pattern in addition to less regular movements. In Chapter 5 we demonstrated daily and weekly *seasonal heteroskedasticity* of FX prices. This seasonality of volatility has been found in intradaily and intraweekly frequencies. In the presence of seasonal heteroskedasticity, autocorrelation coefficients are significantly higher for time lags that are integer multiples of the seasonal period than for other lags. An extended autocorrelation analysis is studied in Chapter 7.

As studied in Chapter 5, the intraweek analysis indicates that the mean absolute returns are much higher over working days than over Saturdays and Sundays, when the market agents are hardly present. The intraday analysis also demonstrates that the mean absolute hourly returns have distinct seasonal patterns. These patterns are clearly correlated to the changing *presence* of main market places of the worldwide FX market. The lowest market presence outside the weekend happens during the lunch hour in Japan (noon break in Japan, night in America and Europe). It is at this time when the minimum of mean absolute hourly returns is found.

Chapter 5 also presents evidence of a strong correlation between market presence and volatility such that the intraday price quotes are positively correlated to volatility when measured with mean absolute hourly returns. Market presence is related to worldwide transaction volume which cannot be observed directly. In the literature, a number of papers present substantial evidence in favor of a positive correlation between returns and volume in financial markets, see the survey of (Karpoff, 1987).

The correlation of market presence and volatility requires us to model and *explain* 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 a high-frequency database (such as the O&A database), which also contains the originating locations of the quoted prices.

In many of the approaches cited in the introduction, in particular in Baillie and Bollerslev (1990) where the seasonality of volatility is modeled by dummy variables, no further explanation of this seasonal pattern is given. We consider it advantageous to try to identify at every moment of the day *which* markets are responsible for the current volatility.

### 6.2.2 Modeling the Volatility Patterns with an Alternative Time Scale and an Activity Variable

Before relating the empirically observed volatility to the market presence, we introduce a model of the price process, which will be used for describing and analyzing the seasonal volatility patterns. A return process with strong intraday and intraweek volatility patterns may not be stationary. Our model for the seasonal volatility fluctuations introduces a new *time scale* such that the transformed data in this new time scale do not possess intraday seasonalities.

The construction of this time scale utilizes two components: the *directing process*,  $\vartheta(t)$ , and a *subordinated* price process generated from the directing process. Let  $x(t)$  be the tick-by-tick financial time series that inherits intraday seasonalities. The directing process,  $\vartheta(t) : \mathbb{R} \rightarrow \mathbb{R}$ , is a mapping from physical time to another predetermined time scale. Here, it is defined such that it contains the intraday seasonal variations.<sup>1</sup>  $\vartheta(t)$ , when used with the subordinated price generating process  $x(t) = x^*[\vartheta(t)]$ , leads to the  $x^*$  process, which has no intraday seasonalities. Although this is not the only possible model to treat the observed seasonality, other traditional deseasonalization techniques are not applicable as the *volatility* is seasonal, not the raw time series.

In the literature, a variety of alternative time scales have been proposed, in different contexts. In the early 1960s, Allais (see, for instance, Allais, 1974) had proposed the concept of *psychological time* to formulate the quantity theory of money. Mandelbrot and Taylor (1967) suggested to cumulate the *transaction volume* to obtain a new time scale which they call the *transaction clock*. Clark (1973) suggested a similar approach. Stock (1988) studied the postwar U.S. GNP and interest rates and proposed a new time scale to model the conditional heteroskedasticity exhibited by these time series. Here we propose to use a new time scale to account for the seasonality.

<sup>1</sup> The  $\vartheta(t)$  process can assume different roles in different filtering environments. If, for instance, the interest is to simply filter out certain holiday effects from the data, then  $\vartheta(t)$  can be defined accordingly. Under such a definition, the transformation will only eliminate the specified holiday effects from the underlying  $x(t)$  process. The  $\vartheta$  type time transformations are not limited to seasonality filtering. They can also be used within other contexts such as the modeling of intrinsic time or transaction clock.

Because the  $\vartheta$ -scale fully accounts for the seasonality of  $x$ ,  $x^*$  has no seasonal volatility patterns. The process  $x^*$  may however have *nonseasonal* volatility patterns; it may be conditionally heteroskedastic. No attempt is made in this chapter to determine its exact nature. The time scale  $\vartheta(t)$  is a strictly monotonic function of physical time  $t$ . Any time interval from  $t_1$  to  $t_2 (> t_1)$  corresponds to a  $\vartheta$ -time interval of the positive size  $\vartheta_2 - \vartheta_1$ . The new *activity* variable  $a$  is defined as the ratio of the interval sizes on the different scales,

$$a_{1,2} \equiv \frac{\vartheta_2 - \vartheta_1}{t_2 - t_1} \quad (6.1)$$

This activity reflects the seasonal volatility patterns. Its relation to other “activity” variables such as market presence or transaction volume was mentioned in Section 6.2.1 and is discussed below.<sup>2</sup>

### 6.2.3 Market Activity and Scaling Law

The volatility-based activity defined by Equation 6.1 can be computed with the empirical *scaling law* (see Chapter 5) for returns, which relates (for  $p = 1$ )  $\langle |\Delta x| \rangle$ , the mean absolute returns over a time interval to the size of this interval,  $\Delta t$ ,


$$E[|r|] = c \Delta t^D \quad (6.2)$$

where  $E$  is the expectation operator,  $c$  is a constant depending of the specific time series.  $D$  is the drift exponent, which determines the scaling properties of the underlying process across different data frequencies. The drift exponent  $D$  is about 0.6 for major FX rates, whereas the pure Gaussian random walk model would imply  $D = 0.5$ . The scaling law expressed in Equation 6.2 holds for all time series studied and for a wide variety of time intervals ranging from 10 min to more than a year.

The scaling law is applied to subsamples in a so-called *intra-week* analysis that allows us to study the *daily* seasonality (open periods of the main markets around the world) as well as the *weekly* seasonality (working days – weekend). Here, we choose a sampling granularity of  $\Delta t = 1$  hr. The week is subdivided into 168 hr 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 is assigned to the corresponding *subsample* with index  $i$ . The 168 subsamples together constitute the full 4-year sample. The sample pattern is independent of bank holidays and daylight saving time. A typical intraday and intra-week pattern across the 168 hr of a typical week is shown in Figure 6.1.

<sup>2</sup> In skipping Saturdays and Sundays, other researchers use an implicit activity model with zero activity on the weekends.





**FIGURE 6.1** Histogram of the average hourly activity (as defined in Equation 6.4 for a statistical week (over 4 years) for the USD-DEM rate.

The scaling law, Equation 6.2, is applied to the  $i^{th}$  hourly subsample instead of the full sample and mathematically transformed to

$$\Delta\vartheta_i = \left( \frac{E[|r_i|]}{c^*} \right)^{1/D} \quad (6.3)$$

From Chapter 5, we know that  $r_i$  can strongly vary for the different hours of a week. The time interval  $\Delta t = 1$  hr (for the hourly sampling) is nevertheless constant. Therefore, it is replaced by the interval  $\Delta\vartheta_i$  on the new time scale  $\vartheta$ . The size of  $\Delta\vartheta_i$  is no longer constant, but reflects the typical volatility of the  $i^{th}$  hour. The constant  $c^*$  is essentially the  $c$  of Equation 6.2, but can differ slightly as it is calibrated by a normalization condition presented later.

The *activity* of the  $i^{th}$  hourly subsample directly follows from Equation 6.1,

$$a_{\text{stat},i} \equiv \frac{1}{\Delta t} \left( \frac{E[|r_i|]}{c^*} \right)^{1/D}, \quad \Delta t = 1 \text{ hr} \quad (6.4)$$

This is the volatility-based activity definition used in the following analysis. The constant  $c^*$  is calibrated to satisfy the following, straightforward normalization condition:

$$\frac{1}{168} \sum_{i=1}^{168} a_{\text{stat},i} = 1 \quad (6.5)$$

#### 6.2.4 Geographical Components of Market Activity

In Figure 6.1, the histogram of the average hourly activity defined by Equation 6.4 is plotted for the USD-DEM rate. Although the activity definition is based only

FOR THE "FUTURE"

on return statistics, the histogram exhibits clear structures where there is very low activity over the weekends and strongly oscillating activity patterns on normal business days. The most active period is the afternoon (GMT) when the European and American markets are open simultaneously. We have varied the  $\Delta t$  granularity of this analysis from 15 min to 4 hr and found no systematic deviations of the resulting activity patterns from the hourly ones. Furthermore, the activity patterns are remarkably stable for each of the 4 years of the total sample. The strong relation between return activity and market presence leads to the explanation of activity as the *sum of geographical components*. Although the FX market is worldwide, the actual transactions are executed and entered in the bookkeeping of particular market centers, the main ones being London, New York, and Tokyo. These centers contribute to the total activity of the market during different market hours that sometimes overlap.

Goodhart and Figliuoli (1992) have explored the geographical nature of the FX market to look for what they call the *island hypothesis*. They studied the possibility that the price bounces back and forth from different centers when special news occurs before finally adjusting to it. Along the same idea, Engle *et al.* (1990), in a study with daily opening and closing USD-JPY prices in the New York and Tokyo markets and a market-specific GARCH model, investigate the interaction between markets. They use the terms *heat wave hypothesis* for a purely market-dependent interaction and *meteor shower hypothesis* for a market-independent autocorrelation. They find empirical evidence in favor of the latter hypothesis. Both studies have not found peculiar behavior for different markets. This encourages us to model the activity with geographical components exhibiting *similar* behavior.

The activity patterns shown in Figure 6.1 and the results reported in Chapter 5 suggest that the worldwide market can be divided into *three* continental components: East Asia, Europe, and America. The grouping of the countries appearing on the Reuters pages in our three components can be found in Table 6.1. This division into three components is quite natural and some empirical evidence supporting it will be presented in Chapter 7.

The model activity of a particular geographical component  $k$  is called  $a_k(t)$ ; the *sum* of the three *additive* component activities is  $a(t)$ :

$$a(t) \equiv \sum_{k=1}^3 a_k(t) \quad (6.6)$$

This total activity should model the intraweekly pattern of the *statistical* activity  $a_{\text{stat},i}$  as closely as possible. Unlike  $a_{\text{stat}}$ , which has relatively complex behavior (see Figure 6.1), the components  $a_k(t)$  should have a simple form, in line with known opening and closing hours and activity peaks of the market centers.

### 6.2.5 A Model of Intra-week Market Activity

Each of the three markets has its activity function  $a_k(t)$ . For modeling this, we use quantitative information on market presence. A statistical analysis of the number

**TABLE 6.1** Definition of the three generic markets.

Grouping of the different countries appearing in the multicontributor pages or record from Reuters according to the three components of the worldwide market.

Index $k$	Component	Countries
1	East Asia	Australia, Hong Kong, India, Indonesia, Japan, South Korea, Malaysia, New Zealand, Singapore
2	Europe	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
3	America	Argentina, Brazil, Canada, Mexico, United States

of price quotes originating from each of the three markets defined by Table 6.1 reveals two aspects on market presence:

- A market has opening times that are *longer* than those of a particular submarket (e.g., an individual bank in one financial center such as Tokyo, Paris, or Chicago). The market opening time is the *union* of the opening times of all relevant institutions of the market.
- Two markets (East Asia and Europe) have a local price quote frequency minimum in the middle of their working day, corresponding to a *noon break*. This local minimum is very pronounced in East Asia and moderate in Europe. In America, there is no minimum around noon. These differences reflect the well-known, different business habits concerning lunch breaks.

Each of the three markets is modeled to have two basic states, either open or closed. The activity does not completely go to zero when the market is closed because it is defined in terms of returns. The activity during the closing hours is modeled to stay on a *small* constant base level  $a_{0,k}$ . During the opening hours, a much stronger, varying, positive activity  $a_{1,k}$  adds to the base level,

$$a(t) \equiv \sum_{k=1}^3 [a_{0,k} + a_{1,k}(t)] \equiv a_0 + \sum_{k=1}^3 a_{1,k}(t) > 0 \quad (6.7)$$

The joint base level  $a_0$  is regarded as one model parameter. There is no need to analyze components  $a_{0,k}$ .

The activity during opening hours,  $a_{1,k}$ , is modeled with a polynomial with smooth transition to the constant behavior of the closing hours. This choice is mathematically convenient because such functions are easily differentiable and

analytically integrable. For parsimony, the number of parameters of this polynomial is kept at a minimum to model the smooth transitions, the lunch break, and the skewness to account for the relative weights of morning and afternoon hours.

In the subsequent analysis, the statistical week is considered from  $t = 0$  on Monday 00:00 to  $t = 168$  hr on Sunday 24:00 (GMT), as shown in Figure 6.1. In order to define the opening and closing conditions of the markets in a convenient form, an *auxiliary* time scale  $T_k$  is introduced. Essentially, it is GMT time; the following market-dependent transformations are only done for technical convenience:

$$T_k \equiv [(t + \Delta t_k) \text{ modulo } (24 \text{ hr})] - \Delta t_k \quad (6.8)$$

where  $\Delta t_k$  has the value of 9 hr for East Asia, 0 for Europe, and  $-5$  hr for America. (The result of the modulo operator is the left-hand side argument *minus* the nearest lower integer multiple of the right-hand side argument.) The *weekend* condition (WEC) also depends on the market:<sup>3</sup>

$$(t + \Delta t_k) \text{ modulo } (168 \text{ hr}) \geq 120 \text{ hr} \quad (6.9)$$

Now the model for an individual market component can be formulated by

$$a_{1,k}(t) \equiv \begin{cases} 0 & \text{if } T_k < o_k \text{ or } T_k > c_k \text{ or (WEC)} \\ a_{\text{open},k}(t) & \text{if } o_k < T_k < c_k \text{ and not (WEC)} \end{cases} \quad (6.10)$$


where  $o_k$  and  $c_k$  are the parameters for the opening and closing hours, respectively. The polynomial function is

$$a_{\text{open},k}(t) \equiv \frac{\omega_k}{\frac{o_k+c_k}{2} - s_k} (T_k - o_k)^2 (T_k - c_k)^2 (T_k - s_k) [(T_k - m_k)^2 + d_k^2] \quad (6.11)$$

where  $\omega_k$  represents the scale factor of the  $k^{\text{th}}$  market,  $s_k$  the skewness of the activity curve,  $m_k$  fixes the place of the relative minimum around the noon break, and  $d_k$  determines the depth of this minimum. The special form of the first factor is chosen to avoid too strong a dependence of the scale factor on  $s_k$ .

In Figure 6.2, the panel on the left illustrates the shape of the geographical seasonality in the European market. The opening and closing times are where the activity level is zero. These parameters are illustrated with “ $o$ ” and “ $c$ ” signs. The seasonality has two peaks with the second peak higher than the former. The relative minimum between the two peaks is the lunch break effect. The location and depth of this relative minimum are controlled by the parameters “ $m$ ” and “ $d$ ” of the last term of Equation 6.11. The activity starts to peak with the opening

<sup>3</sup> The Japanese markets were open on some Saturday mornings according to certain rules in earlier years. These Saturdays, which are noticeable in Figure 6.1, are neglected here, but discussed in Section 6.3.2.



**FIGURE 6.2** The geographical seasonality patterns. The panel on the left illustrates the shape of the geographical seasonality in the European market. The seasonality has two peaks with the second peak higher than the former. The relative minimum between the two peaks is the lunch break effect. In the right panel, the North American geographical seasonality is plotted. It has no lunch break effect.

of the market in the morning, slows down during the lunch break and it peaks in the afternoon again. As the market closing time approaches, the level of activity gradually goes down and reaches zero. In the right panel of Figure 6.2, the North American geographical seasonality is plotted, which has no lunch break effect. The parameter  $s$  controls the asymmetry of the peaks for the European market, whereas in the case of the North American market, it controls the skewness of the overall pattern.

This polynomial model applies to all markets. The European and Asian markets ( $k = 1, 2$ ) have finite  $d_k$  values in the fitting process, but for the American one, the parameter  $d_3$  always diverges to very high values. This reflects the missing noon break in this market, which has already been found in the tick frequency statistics.

The Equation 6.11 for America thus degenerates to a simpler form with no local activity minimum

$$a_{\text{open},3}(t) \equiv \frac{\omega_3}{\frac{o_3+c_3}{2} - s_3} (T_3 - o_3)^2 (T_3 - c_3)^2 (T_3 - s_3) \quad (6.12)$$

Some of the model parameters, the opening and closing times, are already known from the quote frequency statistics. For the other parameters, there are constraints.

To ensure positive activities,  $a_0$  and  $\omega_k$  must be positive and  $s_k$  outside the opening hours,

$$a_0 > 0, \omega_k > 0, s_k \leq o_k \text{ or } s_k \geq c_k \quad (6.13)$$

The parameter  $m_k$  in Equation 6.11 should be within the opening hours as it models the noon break:

$$o_k < m_k < c_k \quad (6.14)$$

The functions  $a_{1,k}(t)$  must be fitted to the results of the statistics,  $a_{\text{stat}}(t)$ , by minimizing the integral of the weighted square deviation of  $a(t)$  from  $a_{\text{stat}}(t)$ . A continuous function  $a_{\text{stat}}(t)$  is not available but rather the hourly series  $a_{\text{stat},i}$  from Equation 6.4. Therefore, the sum over the intraweekly sample is used instead of the integral:

$$\sum_{i=1}^{168} \frac{\left[ a_{\text{stat},i} - a_0 - \sum_{k=1}^3 a_{1,k}(t_i) \right]^2}{\sigma_{\text{error},i}^2} = \min, \quad t_i = \left( i - \frac{1}{2} \right) \text{ hr} \quad (6.15)$$

The hourly intervals are represented by their middle points in this approximation. The least square model has 11 parameters, three  $\omega_k$ 's, three  $s_k$ 's, two  $m_k$ 's, two  $d_k$ 's, and the base activity  $a_0$ . The values of opening,  $o_k$ , and closing,  $c_k$ , are subject to random measurement error originating from the price quote frequency statistics. Therefore, these values are allowed to vary *slightly* for adjusting the fit. The minimization problem of Equation 6.15 is nonlinear in some of the parameters. It can be solved by the Levenberg-Marquardt method (see Press *et al.*, 1986, Section 14.4), but in complex cases a simple genetic search algorithm provides the optimum parameters much more efficiently.

The main American and European markets observe daylight saving time during summer, whereas the main East Asian markets do not. This fact is ignored for the fitting. Only the GMT scale is used. A posterior daylight saving time correction is proposed in Section 6.3.2.

The resulting parameter estimates for four major FX rates and gold (XAU-USD) are presented in Table 6.2 together with the relative weights of the different markets (to be defined in Section 6.3.1). In the top panel of Figure 6.3, the resulting activity model together with the statistical activity for the USD-JPY is shown, and the bottom panel of Figure 6.3 shows the same quantity for the USD-CHF. Figure 6.4 displays the activity model over 48 hr (outside the weekend) with its different components for the same rates.

### 6.2.6 Interpretation of the Activity Modeling Results

The resulting parameters of the activity model and Figures 6.3 and 6.4 confirm the close relation between market presence and intraweekly volatility patterns. The market-specific tick frequency analysis and the activity fitting results compare favorably taking into account the Reuters coverage and the limitations of our model.

**TABLE 6.2** The parameter estimates for the three generic markets.

The parameter estimates for the major FX rates and gold (XAU-USD) with the corresponding market weights. The sum of the market weights is less than 100 percent. The rest is accounted for by the basic activity  $a_0$ . The residual activity  $a_0$ , the scale factor  $\omega$ , and the parameter  $d$ , which determines the depth of the minimum at lunch time, are dimensionless numbers. The  $\omega$  values are a factor of  $10^{-4}$ .

Rate	$a_0$	$k$	Market	Weight	$\omega$	$o$	$c$	$m$	$d$	$s$
USD-DEM	0.03	1	East Asian	24.1%	1.69	-3:32	8:24	3:33	0.97	-3:33
		2	European	38.5%	1.07	5:54	18:39	11:07	2.06	20:21
		3	American	34.1%	12.46	11:24	23:25	—	—	40:44
USD-JPY	0.03	1	East Asian	35.4%	1.40	-4:14	8:43	3:35	1.01	-4:17
		2	European	27.6%	5.37	6:55	16:40	11:02	1.51	17:23
		3	American	33.4%	18.73	11:48	22:50	—	—	34:55
GBP-USD	0.02	1	East Asian	24.3%	1.05	-3:48	8:59	3:40	1.08	-4:02
		2	European	39.1%	0.98	6:00	18:19	11:13	2.85	20:05
		3	American	34.0%	13.88	11:24	23:11	—	—	31:43
USD-CHF	0.01	1	East Asian	22.0%	1.12	-4:00	9:00	3:40	1.06	-4:00
		2	European	45.1%	1.04	5:00	18:00	11:23	2.45	-4:45
		3	American	31.6%	13.71	12:00	24:00	—	—	24:00
XAU-USD	0.02	1	East Asian	9.7%	0.14	-3:43	9:36	4:05	3.17	-4:15
		2	European	54.8%	2.98	5:36	17:19	11:10	1.54	2:42
		3	American	33.8%	354.9	15:21	21:30	—	—	21:32

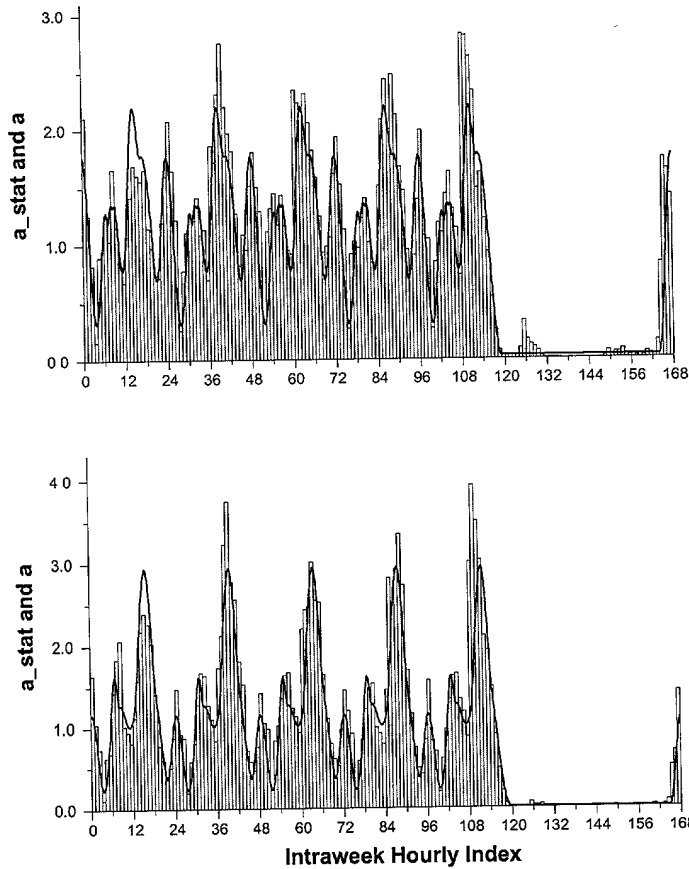
In both cases and for all FX rates, the local minima around noon have the following properties: they are pronounced in East Asia, moderate in Europe, and do not exist in America.

The USD-DEM and the USD-CHF have close parameter values as would be expected with a larger weight for Europe in the case of the USD-CHF, whereas the USD-DEM shows a higher weight for the American market. Gold (XAU-USD) has a very small East Asian market, which extends late because it is mainly traded with Europe. In general, its active trading periods in the individual markets seem to be less extended than for the FX rates. A similar effect is detected with silver. The USD-JPY has a strong East Asian component with a strong overlap with the American market. It is for this rate that the earliest opening of the East Asian market is found. The first example in Figure 6.4 (USD-JPY) has its main market in East Asia. The second example in Figure 6.4 (USD-CHF) has it in Europe, in line with the common sense expectation.

An alternative measure of market activity could also be based on the frequency of price quotes. According to the study in Chapter 5 (Table 5.13), this variable is highly correlated to the volatility. Yet we do not recommend it as an activity measure for two reasons.

1. This number depends on the coverage of the FX market by Reuters and its policy to publish prices on its FFX page. For instance, a new price was

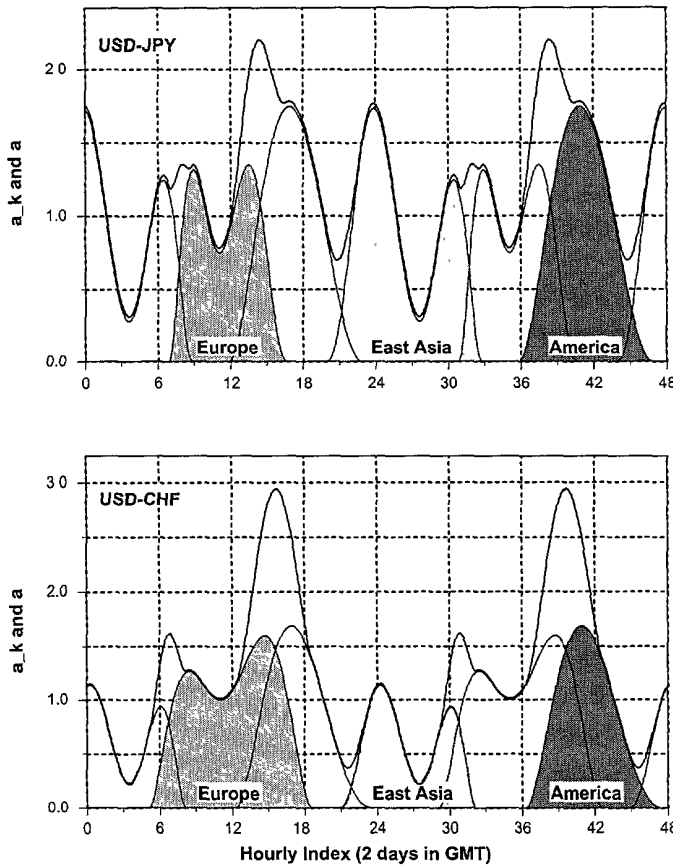




**FIGURE 6.3** The histograms of the average hourly activity for a statistical week (over 4 years) for the USD-JPY (above) and USD-CHF (below) rates and the modeled activity.

shown for a particular rate on this page at maximum one price every 6 sec. Some relevant price revisions were therefore lost because of limitations of the data supplier. Whereas the price revisions depend directly on the data supplier's coverage or policy, the prices are issued by market makers who closely follow the real market value and have many data sources available. Thus published prices are conditioned more by other simultaneously available prices, which do not necessarily appear on this data source.

In order to provide some empirical evidence of this dependence, we compare the hourly shares of the weekly number of price revisions in the 168 hr of the statistical week (see Section 6.2.1) of two different data suppli-



**FIGURE 6.4** The model activity decomposed into the three different continental markets over a period of 48 hr during normal business days for the same rates as in Figure 6.3. The top curve is the sum of  $a_0$  and the three market activities.

ers, Reuters and Knight Ridder,<sup>4</sup> for the same period. The two resulting statistical functions differ substantially. Knight Ridder data are about half as frequent as Reuters data and cover the East Asian markets quite poorly. We measure the difference of the two curves in terms of the root mean squared error (RMSE) of all hourly differences.

We then apply the same approach for a comparison of absolute returns between the two suppliers. We analogously measure the difference between the two resulting curves in terms of the RMSE of the hourly differences.

<sup>4</sup> Since this study was done, Knight Ridder has been integrated with Telerate to Bridge.

"FOREX" IN THE 2000s

The RMSE ratio  $R_{RMSE}$  is defined as follows:

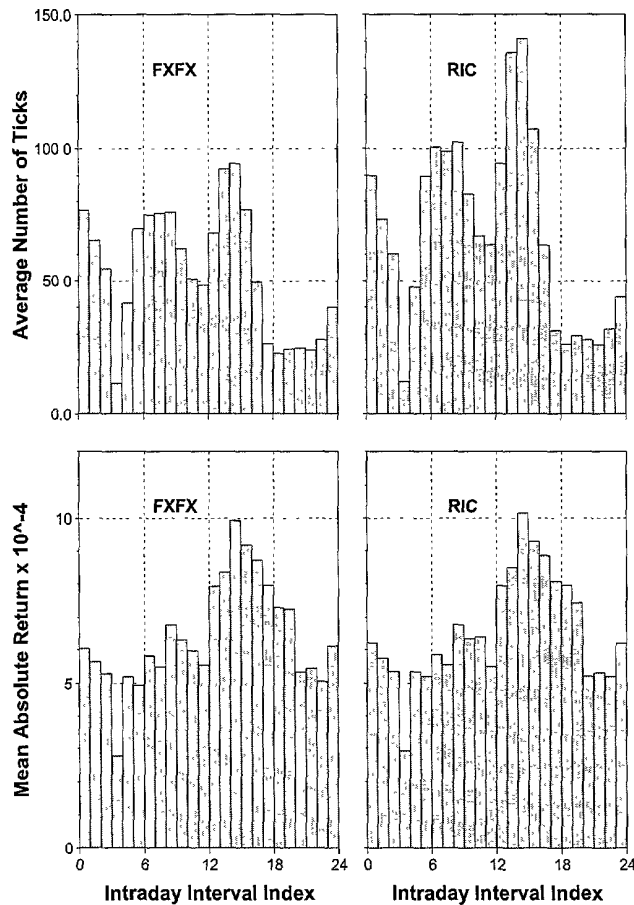
$$R_{RMSE} \equiv \sqrt{\frac{\sum_{i=1}^{24} [(\bar{v}_i^{re} - \bar{v}_i^{kr}) / \bar{v}_i^{re}]^2}{\sum_{i=1}^{24} [(\bar{f}_i^{re} - \bar{f}_i^{kr}) / \bar{f}_i^{re}]^2}} \quad (6.16)$$

where  $\bar{f}_i$  are the mean hourly number of ticks and  $\bar{v}_i$  are the mean hourly absolute returns. The RMSE value here is consistently lower than that for the tick frequency; the ratio is 0.32 for DEM-USD, 0.17 for JPY-USD, 0.20 for USD-GBP, 0.42 for CHF-USD, and 0.52 for XAU-USD. This shows that the volatility is less dependent than the number of price revisions on the data supplier.

Another illustration of this is given in Figure 6.5. We show in these graphs the result of an intraday study of both the tick frequency and the average hourly returns for USD-JPY computed during the same time period on a sample coming from the traditional FFX page of Reuters (left graphs) and another sample coming from the new method Reuters chose to publish its data, the Reuters Instrument Codes (RICs). This new method, being much more suited for computer manipulations, allows the data vendor to transmit much more information and this is very apparent when examining the two upper graphs on the hourly number of ticks. On the other hand, the two lower graphs show little differences because they are computed directly from the prices, which are not governed by the data vendor policy but rather by the market. We use a similar RMSE ratio as in Equation 6.16 and find values around 0.12. This example indicates clearly the problem one is faced with the activity definition. During the same period and for the same market, the activity should be independent of the data source. This is only the case for the hourly absolute returns.

2. Returns are less sensitive than the tick frequency to data holes. The frequency goes to zero if the communication line is broken (there is no good interpolation method for this variable) whereas, with the proper price interpolation, only the variation around the interpolated line for the returns is lost.

The transaction volume, a potential candidate to describe market activity, is not available in hourly frequency. Transaction volume data are available for particular dates through two surveys published by the Federal Reserve Bank of New York (1986 and 1989). Although these surveys are useful to quantify the amount of capital involved, they do not give any indication about intradaily, daily, and weekly changes. We do not propose our activity model as a direct model for the seasonal patterns of transaction volume, but suggest its usefulness in future research.



**FIGURE 6.5** The comparison of the tick activity (upper graphs) and the hourly absolute return (lower graphs) for two data sources. The old Reuters FFX page and the new Reuters Instrument Code (RIC) data. The comparison is conducted for the USD-JPY from October 25, 1993 to March 18, 1995.

### 6.3 A NEW BUSINESS TIME SCALE ( $\vartheta$ -SCALE)

#### 6.3.1 Definition of the $\vartheta$ -Scale

In Section 6.2.2, the time scale  $\vartheta$  was introduced to model the seasonal, intradaily and intraweekly aspect of heteroskedasticity. In Equation 6.1, the activity variable has been defined as the “speed” of  $\vartheta$  against the physical time  $t$ . The continuous activity function  $a(t)$  of Equation 6.7, developed in the previous sections, allows us to define  $\vartheta$  as its time integral,

$$\vartheta \equiv \vartheta(t) \equiv \int_{t_0}^t a(t') dt' \equiv a_0 (t - t_0) + \sum_{k=1}^3 \vartheta_k(t) \quad (6.17)$$

The starting date  $t_0$  chosen for the  $\vartheta$ -scale is arbitrary. The activity is always positive, so its integral  $\vartheta(t)$  is a monotonically increasing function. The  $\vartheta_k$  represents in fact the business time scale of the  $k^{\text{th}}$  market and is defined as

$$\vartheta_k(t) \equiv \int_{t_0}^t a_{1,k}(t') dt' \quad (6.18)$$

This quantity is informative in itself and can be used to model intramarket behavior. Because of the regular weekly pattern of  $a$ ,  $\vartheta$  is *predictable* according to Equation 6.17; it may be computed also for the future. Due to *normalization* (see Equation 6.5),  $\vartheta$ -time can be measured in the *same units* as physical time (e.g., hours, days, weeks); one full week in  $\vartheta$ -time corresponds to one week in physical time.

The *relative weight*,  $W_k$ , of each market component can be defined with the help of the integral  $\vartheta_k$  over a full week:


$$W_k = \frac{\vartheta_k(t + 1 \text{ week}) - \vartheta_k(t)}{\vartheta(t + 1 \text{ week}) - \vartheta(t)} = \frac{\vartheta_k(t + 1 \text{ week}) - \vartheta_k(t)}{1 \text{ week}} \quad (6.19)$$

This is the share of the  $k^{\text{th}}$  market in the  $\vartheta$  interval of one week. In Table 6.2, the relative weights of each component, as given by Equation 6.19, are presented together with the fitted parameters. These weights are in fact interesting pieces of information about the market shares of the components defined in Section 6.2.4 and Table 6.1. They are in line with the results of the market surveys regularly made by the (Bank for International Settlements, 1990, 1993, 1995).

The  $\vartheta$ -scale contracts periods of low activity and expands period of high activity. This is clearly seen on Figure 6.6 where the mapping function between  $\vartheta$ -time and physical time is shown for USD-DEM over a week. Because the  $\vartheta$ -time is normalized to physical time over 4 years (see the next section), the two scales almost coincide after a week but not exactly ( $\vartheta$ -scale is slightly above 168 hr), because we have chosen the week of September 9 to October 1, 1995, where there was no market holiday. This figure shows that during the weekend,  $\vartheta$ -time flows very slowly, compensating for the low activity during this period in physical time.

### 6.3.2 Adjustments of the $\vartheta$ -Scale Definition

The  $\vartheta$ -scale defined in Section 6.3.1 reflects a rigid intraweekly pattern of expected market activity. However, there is more relevant information about the activity due to information on business holidays, daylight saving times, and scheduled events in general. In practice, for volatility forecasts, it is desirable to account for this information in the construction of the  $\vartheta$ -scale. Such adjustments are carried out in Equation 6.17 by recalibrating the factor  $c^*$  over the whole sample.




**FIGURE 6.6** The time mapping function between physical time and  $\vartheta$  time. The week chosen to draw this mapping function is a week with no market holidays (September 25 to October 1, 1995). The thin line represents the flow of physical time.

It is difficult to take into account the different *holidays* of *each* market accurately.<sup>5</sup> In the framework of the three markets of Table 6.1, our approach is an approximate solution. A holiday is considered if it is common to a large part of one of the three markets of the model. On such holidays, the activity  $a_{1,k}$  is set to zero for this market. The holiday is treated like a weekend day in Equation 6.10.

In some countries, there are *half-day holidays*. Their treatment would require the splitting of the daily activity functions into morning and afternoon parts. This splitting could also be used to model the few Saturday mornings in Japan (until 1989) when the banks were open. These modifications have not been made as they are beyond our objective of modeling the *main* features of the FX activity patterns.

The *daylight saving time* observed in two of the markets, Europe and America, has an influence on the activity pattern and thus on  $\vartheta$ . The presence of local markets depends on local time rather than on GMT. One way to deal with this is to convert the time constants of Table 6.2 from GMT to a typical local time scale of the

<sup>5</sup> Future holidays are not always known in advance as, for instance, the Islamic holidays. Thus,  $\vartheta$  might no longer be predictable in those special cases.



**FIGURE 6.7** The hourly returns for USD-DEM from June 3, 1996, 00:00:00 to September 11, 1996, 00:00:00 are plotted using the physical time scale and the  $\vartheta$ -time scale. Note also the extreme events that are clearly visible on both graphs.

market. This conversion yields different results for the local times in summer and in winter. The time constants are fixed to the *mean* of the summer and winter conversion results, reflecting the fact that the sample used in the activity fitting is composed of approximately half summer and half winter. The computation of the activity and  $\vartheta$  is then based on Equation 6.10 with these *local* time constants. A better algorithm, which takes into account the difference between summer and winter local market time and which allows a dynamic adaptation to changes in the activity pattern indicates substantial improvement (Breymann, 2000).

So far, volatility patterns with periods of more than one week have been neglected. Yet there may be patterns with longer periods caused by month-end effects, by the monthly or quarterly releases of certain important figures such as the American trade or unemployment figures, and by yearly effects. Moreover, there are long-term changes such as the overall volatility increase over the past 15 years as shown in Chapter 5. None of these effects has been found to be significant in a 4-year sample we studied.

Figure 6.7 illustrates the effect of the time transformation with the hourly returns of USD-DEM over 3 months both in physical and in  $\vartheta$ -time. It is easy to

2025 RELEASE UNDER E.O. 14176



**TABLE 6.3** Quality test of the  $\vartheta$ -scale.

Test for the quality of the  $\vartheta$ -scale as calculated in Equation 6.21. This ratio illustrates the reduction of intraweekly volatility fluctuations when using the  $\vartheta$ -scale.

	USD-DEM	USD-JPY	GBP-USD	USD-CHF	XAU-USD
Volatility ratio	0.28	0.29	0.25	0.29	0.25

see that the quiet periods during the weekends are in the upper graph in physical time. They give the sense of periodicity. In the lower graph, where hourly returns are computed in  $\vartheta$ -hours, the seasonality is removed and the picture resembles much more those made with weekly or daily data (omitting weekends). Another remarkable feature of these graphs is the number of large movements. During this period, the USD-DEM experienced price changes as high as 1.5% in an hour.

### 6.3.3 A Ratio Test for the $\vartheta$ -Scale Quality

There are various ways to measure the quality of a  $\vartheta$ -time scale. Because the goal of such a scale is to remove the daily and weekly seasonality of volatility, it is natural to test the extent to which this has been achieved. Here we define a quantitative test that allows discrimination between various possible business time scales.

The absolute returns on an intraweekly sample as described in Section 6.2.3 are first computed on the physical time scale. We define the size of the weekly fluctuations of mean volatility:

$$F_{v(t)} = \left[ \frac{1}{N} \sum_{i=1}^N \left( \frac{\sum_{j=1}^m |r_j(t_i)|}{m} - \frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^m |r_j(t_i)|}{m} \right)^2 \right]^{1/2} \quad (6.20)$$

where  $i$  is the index of the hourly interval in the statistical week and  $N = 168$  the total number of these intervals. Absolute returns are observed and averaged over  $m$  weeks with index  $j$ , for each hour of the statistical week. The fluctuations, which are large when analyzed in physical time  $t$ , should be strongly reduced when analyzed in  $\vartheta$ -time. For analyzing the fluctuations in  $\vartheta$ -time, the sampling over one full week is again divided into 168 intervals. Instead of being equally spaced in physical time, they are now equally spaced in  $\vartheta$ -time. This condition can be formally written as  $\vartheta(t_{i+1}) - \vartheta(t_i) = 1$  hr, where the hour is now measured on the  $\vartheta$ -scale. The sequence  $t_i$  that fulfills this condition is computed by numerical inversion of the  $\vartheta(t)$  function on one week. The volatility ratio is defined by

$$F_{v(\vartheta)} / F_{v(t)} \quad (6.21)$$

where  $F_{v(\vartheta)}$  and  $F_{v(t)}$  measure the deseasonalized and raw volatility fluctuations. This ratio measures the quality of the extent to which the  $\vartheta$  scale successfully eliminates the seasonal fluctuations of the volatility.

In Table 6.3, the resulting ratio is between 0.25 and 0.29 for all rates indicating the quality of the  $\vartheta$ -scale. For a perfect  $\vartheta$ -scale, the measure tends to zero, and

for physical time, the measure is one. Any other  $\vartheta$ -scale derivation can also be measured the same way, the one with the lowest ratio being the best intraday deseasonalization method. In the next chapter, we will utilize the  $\vartheta$ -scale in analyzing the autocorrelation function of absolute returns.

#### 6.4 FILTERING INTRADAY SEASONALITIES WITH WAVELETS

The previous sections show that the practical estimation and extraction of the intraday periodic component of the return volatility is feasible. The literature also demonstrated that such extraction of the seasonal volatility component is indispensable for meaningful intraday studies. Earlier studies have shown that strong intraday seasonalities may induce distortions in the estimation of volatility models and are also the dominant source for the underlying misspecifications as studied in (Guillaume *et al.*, 1994; Andersen and Bollerslev, 1997b). Besides, Section 7.3 reveals how such a periodic component pulls the calculated autocorrelations down, giving the impression that there is no persistence other than particular periodicities.


To illustrate the impact of seasonalities, Gençay *et al.* (2001a) consider the following AR(1) process with a periodic component:

$$y_t = \alpha + \beta y_{t-1} + \sum_{i=1}^4 3.0 S_{it} + \epsilon_t \quad t = 1 \dots T \quad (6.22)$$

where  $S_{it} = \sin(\frac{2*\pi}{P_i}t) + \eta v_{it}$ ,  $\alpha = 0.0$ ,  $y_0 = 1.0$ ,  $\beta = 0.99$ , and  $T = 1000$ . Periodic components are  $P_1 = 3$ ,  $P_2 = 4$ ,  $P_3 = 5$ , and  $P_4 = 6$  so that the process has 3, 4, 5, and 6 period stochastic seasonality. The random variables  $\epsilon_t$  and  $v_{it}$  are identically and independently distributed disturbance terms with zero mean. The signal-to-noise ratio,  $\eta$ , in each seasonal component is set to 0.30.

Figure 6.8 presents the autocorrelation of the simulated AR(1) process with and without the periodic components. The autocorrelation of the AR(1) process without seasonality (excluding  $\sum 3.0 S_{it}$  from the simulated process) starts from a value of 0.95 and decays hyperbolically as expected. However, the autocorrelation of the AR(1) process with the seasonality indicates the existence of a periodic component. The underlying persistence of the AR(1) process in the absence of the seasonality component is entirely obscured by these periodic components. An obvious route is to filter out the underlying seasonalities from the data. A simple method for extracting intraday seasonality that is free of model selection parameters is proposed by Gençay *et al.* (2001a). The proposed method is based on a wavelet<sup>6</sup> multiscaling approach which decomposes the data into its low and high-frequency components through the application of a nondecimated discrete wavelet transform. In Figure 6.8, the solid line is the autocorrelation of the nonseasonal AR(1) dynamics and the dotted line is the autocorrelation of the deseasonalized series with the method proposed in Gençay *et al.* (2001a). As

<sup>6</sup> An introduction to wavelets can be found in a book by Gençay *et al.* (2001b).




**FIGURE 6.8** Sample autocorrelations for the simulated AR(1) process (straight line), AR(1) plus seasonality process (dot-dashed line), and wavelet transformation of the AR(1) plus seasonality process (dotted straight line).


Figure 6.8 demonstrates, wavelet methodology successfully uncovers the nonseasonal dynamics without imposing any spurious persistence into the filtered series.

With this method, Gençay *et al.* (2001a) study two currencies, namely the 5-min Deutschemark – U.S. Dollar (USD-DEM) and Japanese Yen – U.S. Dollar (USD-JPY) price series for the period from October 1, 1992, to September 29, 1993. This data set is also known as the HFDF-I data set. Figure 6.9 presents autocorrelations of the 5-min absolute return series. This shows that the intradaily absolute returns exhibit strong intraday seasonalities. This phenomenon is well-known and reported extensively in the literature; (see for example, Dacorogna *et al.*, 1993; Andersen and Bollerslev, 1997a).

For a long memory process (see Hosking, 1996), the autocovariance function at lag  $k$  satisfies  $\gamma(k) \sim \lambda k^{-\alpha}$  where  $\lambda$  is the scaling parameter and  $\alpha \in [0, 1]$ . A leading example is the fractionally integrated process for which  $\alpha = 1 - 2d$  and  $d$  is the order of fractional integration. In Andersen and Bollerslev (1997a), the fractional order of integration is estimated as  $d = 0.36$  for the same USD-DEM series utilized in this example. Andersen *et al.* (2001) calculate six  $d$  estimates from various volatility measures for the USD-DEM and USD-JPY series. These six  $d$  estimates vary from 0.346 to 0.448. In this example, the fractional integration parameter is set  $d = 0.4$  to represent the average of these six estimates. Figure 6.10 presents the autocorrelograms of the filtered 5-min absolute returns along with the estimated autocorrelogram of a long memory process with  $d = 0.4$ . These



**FIGURE 6.9** Sample autocorrelations for the USD-DEM and USD-JPY for the 5-min absolute returns of (a) USD-DEM absolute returns and (b) USD-JPY absolute returns from October 1, 1992, through September 29, 1993.



**FIGURE 6.10** Sample autocorrelations for the wavelet filtered 5-min absolute returns of (a) USD-DEM and (b) USD-JPY from October 1, 1992 through September 29, 1993. The dotted line is the autocorrelogram for the estimated hyperbolic decay rate for  $d = 0.40$ —that is,  $k^{-2.0}$  where  $k$  is the number of lags.

findings indicate that the wavelet method is more successful in filtering out intraday seasonalities relative to the method presented in Andersen and Bollerslev (1997a). The persistence of volatility in further lags is also much smaller in Gençay *et al.* (2001a) relative to the Andersen and Bollerslev (1997a). However, the seasonality filters of both Gençay *et al.* (2001a) and Andersen and Bollerslev (1997a) suffer from the fact that the decay of the volatility persistence is slow in the immediate lags relative to the method of Dacorogna *et al.* (1993).

2025 RELEASE UNDER E.O. 14176

# 7

---

## REALIZED VOLATILITY DYNAMICS

### 7.1 INTRODUCTION

High-frequency returns no longer exhibit the seasonal behavior of volatility when investigated in deseasonalized form. Therefore, well-known stylized facts start to be visible in the deseasonalized returns and the corresponding absolute returns. Deseasonalization can be achieved by taking returns regularly spaced in  $\vartheta$ -time. Absolute returns are just one form of realized volatility whose general definition is given by Equation 3.8.

Realized volatility has a considerable statistical error, which can be reduced by taking returns over short time intervals. This leads to a high number of observations within a given sample.<sup>1</sup> Unfortunately, the choice of a small return interval also leads to a bias caused by microstructure effects. This bias is explained in Section 5.5.3 as a consequence of biased quoting, which leads to a bouncing effect of quotes within a range related to the bid-ask spread. In Section 5.5.3, the bias is treated as a component of the measurement error. In Section 7.2, we study the bias empirically and propose a simple bias correction method that applies to the bias caused by any microstructural effect, not only bid-ask bouncing. Bias-corrected realized volatility has a smaller error than the error attainable without correction.

---

<sup>1</sup> Using overlapping returns is also helpful, as explained in Section 3.2.8.

After appropriately defining realized volatility, we can analyze its dynamical behavior through different statistical methods. The fundamental properties of the volatility dynamics are the conditional heteroskedasticity (also called the volatility clustering) and the long memory of the autocorrelation of volatility.<sup>2</sup> In this chapter, we also examine the asymmetry of information flow between volatilities computed from returns measured at different frequencies which is a typical property to study with high-frequency data. Financial markets are made of traders with different trading horizons. In the heart of the trading mechanisms are the market makers. At the next level up are the intraday traders who carry out trades only within a given trading day but do not carry overnight positions. Then there are day traders who may carry positions overnight, short-term traders and long-term traders. Each of these classes of traders may have their own trading tool sets consistent with their trading horizon and may possess a homogeneous appearance within their own classes. Overall, it is the sum of the activities of all traders for all horizons that generates the market prices. Therefore, market activity would not exhibit homogeneous behavior, but the underlying dynamics would be heterogeneous with each trading horizon (trader class) dynamically providing feedback across all trader classes. Figure 7.1 illustrates such a heterogeneous market where a low-frequency shock to the system penetrates through all layers reaching the market maker in the middle. The impact of these low-frequency shocks penetrates the entire market. The high-frequency shocks, however, would be short lived and may have no impact outside their boundaries. We will study this heterogeneity-driven asymmetry in this chapter.


This book utilizes the deseasonalization method explained in Chapter 6, and Dacorogna *et al.* (1993), but a flurry of alternative ways of treating the seasonality have also been proposed: the time-of-day dummy variables, Baillie and Bollerslev (1990); a renormalization of the returns by the seasonal volatility, Taylor and Xu (1997); the flexible Fourier framework to model the seasonal pattern, Andersen and Bollerslev (1997b); time deformation with tick frequency, Pecan *et al.* (1995); Baestaens and Van den Bergh (1995); the use of cubic splines, Engle and Russell (1997); models that include both systematic components and stochastic seasonal components, Beltratti and Morana (1998); and the wavelet multiresolution method of Gençay *et al.* (2001a) in Section 6.4.

## 7.2 THE BIAS OF REALIZED VOLATILITY AND ITS CORRECTION

Realized volatility plays a key role both for the exploration of stylized facts and for practical applications such as market risk assessment. When computing it,

<sup>2</sup> This clustering property was first noted in Mandelbrot (1963) in his study of cotton prices and the long memory in Mandelbrot (1971). These findings remained dormant until the early 1980s for the volatility clustering until Engle (1982) and Bollerslev (1986) proposed the ARCH and GARCH processes. In the early 1990s, a comprehensive study of the long memory properties of the financial markets had started.





**FIGURE 7.1** Financial markets are made of traders with different trading horizons. In the heart of the trading mechanisms are the market makers. A next level up are the intraday traders who carry out trades only within a given trading day. Then there are day traders who may carry positions overnight, short-term traders and long-term traders. Each of these classes of traders may have their own trading tool sets and may possess a homogeneous appearance within their own classes. Overall, it is the sum of the activities of all traders for all horizons that generates the market prices. Therefore, market activity is heterogeneous with each trading horizon dynamically providing feedback across the distributions of trading classes.

using Equation 3.8, we can take advantage of high-frequency data by choosing a short time interval  $\Delta t$  of the analyzed returns. This leads to a large number of observations within a given sample and thus a low stochastic error. At the same time, it leads to a considerable bias in most cases.

In the following bias study, Equation 3.8 is considered in the following form:

$$v(t_i) = v(\Delta t, n, 2; t_i) = \left\{ \frac{1}{n} \sum_{j=1}^n [r(\Delta t; t_{i-n+j})]^2 \right\}^{1/2} \quad (7.1)$$

The choice of the exponent  $p = 2$  has some advantages here. In Section 5.5, we found that the empirical drift exponent of  $v$  is close to the Gaussian value 0.5 if  $v$  is defined with an exponent  $p = 2$ . Assuming such a scaling behavior and a fixed

T O S C O T H R E S

sample of size  $T = n\Delta t$ ,  $v^2$  has an expectation independent of  $\Delta t$ :

$$E[v^2(n\Delta t, 1, 2; t_i)] = n E[v^2(\Delta t, n, 2; t_i)] = \sum_{j=1}^n [r(\Delta t; t_{i-n+j})]^2 \quad (7.2)$$

Thus  $v^2$  can be empirically estimated as the sum of all squared returns within  $T$ , irrespective of the size of  $\Delta t$ . Moreover, the time scale can be changed, such as from  $\vartheta$ -time to physical time, and the return intervals can be of irregular size. This implies that the estimator is also immune to data gaps within the full sample. If prices are interpolated, previous-tick interpolation (see Equation 3.1) should be used here, because linear interpolation leads to an underestimation of volatility. With all the mentioned modifications, the sum of squared returns remains an estimator for  $v^2$ , as long as all the return intervals exactly cover the full sample  $T$ . These nice properties may have led Andersen *et al.* (2000) to choose the name “realized volatility” for the sum of squared returns, as on the right-hand side of Equation 7.2.


The empirically found bias violates Equation 7.2, especially if  $\Delta t$  is very small. The deviation of the empirical behavior from Equation 7.2 provides a measure of the bias. We choose a large enough time interval  $\Delta t_{\text{ref}} = q\Delta t$  as the bias-free reference case to judge the bias of smaller intervals  $\Delta t$ . In practice, a good choice of  $\Delta t_{\text{ref}}$  is between few hours and 1 working day. We define the bias factor  $B(t_i)$ :

$$B(t_i) = \frac{\sqrt{q} v(\Delta t, m, q, 2; t_i)}{v(\Delta t_{\text{ref}}, m, 2; t_i)} = \sqrt{\frac{\sum_{j=1}^{mq} [r(\Delta t; t_{i-mq+j})]^2}{\sum_{j=1}^m [r(\Delta t_{\text{ref}}; t_{i-mq+jq})]^2}} \quad (7.3)$$

where  $m$  is the number of analyzed reference intervals of size  $\Delta t_{\text{ref}}$ , and  $q = \Delta t_{\text{ref}}/\Delta t$  is an integer number. If the scaling assumption of Equation 7.2 is true,  $B(t_i)$  converges to 1 for large samples (i.e., large  $m$  and  $q$ ). The bias can be measured in terms of the deviation of  $B(t_i)$  from 1.

In Figure 7.2, the bias factor  $B(t_i)$  is plotted versus time, for two different markets: the FX rate USD-CHF and the equity index Nikkei-225. The time scale in both cases is a business time: the 49 weekend hr from Friday 8 p.m. GMT to Sunday 9 p.m. GMT are compressed to the equivalent of only 1 hr outside the weekend. The results do not strongly depend on this choice. Similar bias behaviors are obtained when the analysis is done in  $\vartheta$ -time or physical time. The reference time interval is  $\Delta t_{\text{ref}} = 1$  working day. The investigated return intervals  $\Delta t$  are much shorter and vary between 2 min ( $q = 720$ ) and 1 hr ( $q = 24$ ). The number  $m = 260$  of reference intervals is chosen high enough to limit the stochastic error of  $v(\Delta t_{\text{ref}}, m, 2; t_i)$ . This means a bias measurement on a moving sample of about 1 year  $\approx 260$  working days.

The bias factor distinctly deviates from 1 in Figure 7.2, especially for small values of  $\Delta t$  such as 2 min and 5 min. For  $\Delta t = 1$  hr, the bias is still visible but



**FIGURE 7.2** Bias factors plotted versus time, for the FX rate USD-JPY (upper panel) and the Japanese equity index Nikkei-225 (lower panel). Deviations from 1 indicate a bias in realized volatility. The bias factor is the ratio of two mean realized volatilities over the same sample (see Equation 7.3). The investigated return measurement intervals  $\Delta t$  are as follows. Bold curves:  $\Delta t = 1$  hr; middle curves:  $\Delta t = 5$  min; thin curves:  $\Delta t = 2$  min.

can be neglected more easily. Surprisingly, the biases have different signs. The bias of the foreign exchange (FX) rate is positive, whereas that of the equity index is negative ( $B(t_i) < 1$ ). The bias can be explained by microstructure effects, but these are obviously different for different markets. The microstructure effects of FX rates were discussed in Chapter 5, in particular the negative autocorrelation due to a bouncing effect within the bid-ask spread (Section 5.2.1). The bias due to this effect can be modeled as in Section 5.5.3 and in Corsi *et al.* (2001), where the influence of data gaps on the bias is also analyzed. There is ongoing research aiming at refined versions of this bias model. The negative bias of the equity index has to be explained differently. An equity index is a weighted average of some equity prices. Some of the individual equities play a leading role in price adjustments and establish small trends that the other equities follow. This mechanism causes a short-term (few minutes) positive autocorrelation of the index returns and eventually a negative bias of realized volatility when a very short

interval  $\Delta t$  is chosen. The bias factors moderately fluctuate over time, but there are no dramatic shifts. The overall levels are maintained even over the 10-year sample of Figure 7.2 (upper panel).

The bias can be avoided either by taking large return intervals  $\Delta t$  (with the disadvantage of large stochastic errors) or by introducing a *bias correction* for small intervals  $\Delta t$ . Eliminating the bias seems to be a demanding task requiring a model of the microstructure effects. Section 5.5.3 has such a model for FX rates, but other markets such as equity indices need other models.

Instead of developing bias models for each market, we suggest a simple bias correction method that needs no explicit model and only relies on the assumption that the bias-generating mechanism is much more stable over time than the volatility itself. The limited size of bias fluctuations in Figure 7.2 justifies this assumption. The bias correction is simple. Each realized volatility observation is divided by the bias factor as measured in the past:

$$v_{\text{corr}}(\Delta t, n, 2; t_i) = \frac{v(\Delta t, n, 2; t_i)}{B(t_i)} \quad (7.4)$$


where  $B(t_i)$  is defined by Equation 7.3. This bias correction can be computed in real time, because it is based on information fully available at time  $t_i$ . Some variations of Equation 7.4 are possible, as suggested by Corsi *et al.* (2001). The bias correction factor can be computed by moving average operators as explained in Section 3.3 instead of the sums of Equation 7.3.

Figure 7.3 probes the success of the simple bias correction. The bias factor  $B_{\text{corr}}$  of the already bias-corrected realized volatility can be measured in the same way as the bias of the uncorrected volatility (Equation 7.3):

$$B_{\text{corr}}(t_i) = \frac{\sqrt{q} v_{\text{corr}}(\Delta t, m, q, 2; t_i)}{v(\Delta t_{\text{ref}}, m, 2; t_i)} \quad (7.5)$$

A perfect bias correction implies  $B_{\text{corr}}(t_i) \equiv 1$ . However, the bias correction is not perfect. Both the bias correction and its measurement in Equation 7.5 rely on a quantity  $v(\Delta t_{\text{ref}}, m, 2; t_i)$ , which has a stochastic error. These imperfections are visible in the form of fluctuations of  $B_{\text{corr}}$  about 1 in Figure 7.3. Figure 7.2 and Figure 7.3 are based on the same samples and parameters and can directly be compared.  $B_{\text{corr}}$  in Figure 7.3 is much closer to 1 than  $B$  in Figure 7.2, in all cases. This fact demonstrates a successful bias correction for both markets, FX and the equity index.

In spite of the success of Equation 7.4 as shown in Figure 7.3, the simple bias correction has some shortcomings, one of them being the multiplicative nature of the formula. Realized volatility values are corrected by a slowly varying correction factor, irrespective of the current volatility level. One can argue that an additive or nonlinear correction of realized volatility would reflect reality better than the multiplicative correction. (An additive correction may lead to impossible negative volatility values, though.) A fair judgment may be as follows. Equation 7.4



**FIGURE 7.3** Bias factors plotted versus time, for the FX rate USD-JPY (upper panel) and the Japanese equity index Nikkei-225 (lower panel), computed by Equation 7.5. The investigated realized volatility values have already been bias-corrected by Equation 7.4, so the small deviations from 1 indicate imperfections of the bias correction. The investigated return measurement intervals  $\Delta t$  are as follows. Bold curves:  $\Delta t = 1$  hr; middle curves:  $\Delta t = 5$  min; thin curves:  $\Delta t = 2$  min. The same scaling as in Figure 7.2 is used.

succeeds in largely reducing the bias and is thus better than no bias correction. As soon as an appropriate model of the bias-generating process for a particular market exists, the corresponding bias-correction method will be clearly superior to Equation 7.4.

Bias correction is a means to compute realized volatility with smaller intervals  $\Delta t$  and, for a given sample of size  $T = n\Delta t$ , a smaller stochastic error. Unfortunately, the bias correction introduces an additional stochastic error due to the factor  $v(\Delta t_{\text{ref}}, m, 2; t_i)$  in Equation 7.3. Corsi *et al.* (2001) show that a bias-corrected volatility with reasonable parameters has a total error that is still distinctly smaller than the error of uncorrected volatility. The following rough calculation also shows this. Uncorrected volatility requires a rather large  $\Delta t$  of about 1 hr (with  $q = 24$ ) to keep the bias at bay. The stochastic error is proportional to  $\sqrt{1/24}$ , and a bias of roughly of the same size adds to the error. Bias-corrected volatility can have a

small  $\Delta t = 5$  min ( $q = 288$ ). The stochastic error is proportional to  $\sqrt{1/288}$ , but the factor  $v(\Delta t_{\text{ref}}, m, 2; t_i)$  with  $m = 260$  leads to another stochastic error component proportional to  $\sqrt{1/260}$ . Both error components together are proportional to  $\sqrt{1/288 + 1/260} \approx \sqrt{1/137}$ . This is distinctly smaller than the value without bias correction,  $\sqrt{1/24}$  (where the bias makes the error even larger).

So far, the bias discussion has been restricted to realized volatility with an exponent  $p = 2$  in Equation 3.8. When choosing another exponent (such as  $p = 1$ , a good choice for many following studies), the bias discussion becomes more complicated. The scaling behavior deviates from Gaussian scaling, as seen in Section 5.5, and data gaps have a stronger influence on the bias than in the case  $p = 2$ . For exponents other than 2, a bias correction with a formula such as Equation 7.4 is less successful, and more research is needed. The technique of bias correction is rather new and will be improved by ongoing research. The realized volatility studies of the following sections are older and do not contain any bias correction. However, the choice of very short return intervals (such as 5 min) has been avoided, so the size of the bias is limited.

### 7.3 CONDITIONAL HETEROSKEDASTICITY

#### 7.3.1 Autocorrelation of Volatility in $\vartheta$ -Time


This section analyzes the autocorrelations of returns and realized volatility in physical and  $\vartheta$ -time.<sup>3</sup> The study utilizes a 20-min frequency instead of an hourly one. We did not take smaller intervals than 20 min in order to avoid a strong bias, as explained in Section 7.2. The autocorrelation function of the USD-DEM is shown in Figure 7.4 for up to 720 lags. The confidence intervals in Figure 7.4 refer to 95% confidence for a Gaussian random process around the sample mean. Because the distributions of returns and volatility are not Gaussian, the confidence intervals are provided as a reference rather than for exact statistical significance.

In Figure 7.4, the autocorrelation function of volatility has a distinct structure, which is far beyond the confidence intervals. For lags of any integer number of days, clear peaks are found. These peaks indicate the *daily* seasonality. The *weekly* seasonality is highly visible in the form of high autocorrelation for lags around 1 week and low 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 8 hr, corresponding to a frequency *three* times the daily frequency. Our world market model with *three* continental markets is confirmed by this observation. Apart from these seasonal peaks there must be a positive component of the autocorrelation that declines with increasing lag. In Figure 7.4, this component cannot be observed as it is overshadowed by seasonality.

The autocorrelations of returns, unlike those of volatility (absolute returns), are close to zero and within the confidence intervals for most of the lags. The

<sup>3</sup> Absolute returns are studied here.

ACCEPTED MANUSCRIPT




**FIGURE 7.4** The autocorrelation function of USD-DEM returns and volatility (absolute returns). The data sampling is in 20-min frequency in *physical time* for lags up to 10 days. The 95% confidence interval is for a Gaussian random process. The sampling period is from March 3, 1986, to March 3, 1990.

squared returns, instead of absolute returns, may also be used as a proxy for the underlying volatility. Autocorrelations of square returns also exhibit similar seasonality peaks as those of absolute returns, but are less pronounced. It is well known that the theoretical autocorrelation of squared returns is meaningful only if the *kurtosis* of the return process is finite, which is not guaranteed for currency returns.

A similar autocorrelation analysis is also carried out with the  $\vartheta$ -time scale instead of the physical time  $t$ , and it is presented in Figure 7.5. There are no large seasonal peaks in the volatility autocorrelations of the  $\vartheta$ -time. This is due to the fact that the  $\vartheta$ -scale is constructed to eliminate the intraday seasonality. The autocorrelation of volatility is significantly positive and declines at an hyperbolic





**FIGURE 7.5** The autocorrelation function of the USD-DEM returns and the absolute returns at 20-min data frequency in  $\vartheta$ -time. The number of lags is up to 10  $\vartheta$  days. The first lag is marked by an empty circle. The exponential decay is shown with a dashed line. The hyperbolic decay fits best to the autocorrelation function of the absolute returns. The figure on the right is the same autocorrelation function for the absolute returns extended to a much larger number of lags with the superimposition of the hyperbolic decay.

rate. This behavior can be explained by the presence of a long memory process in the underlying data-generating process of returns. The rate of decline in the autocorrelation is, however, slower than an exponential decline, which would be expected for a low-order GARCH process, Bollerslev (1986).

The autocorrelation function of volatility (Figure 7.5) is not completely free of seasonalities. A narrow peak can be identified at a lag of 1 week. This peak might be due to the day of the week effects. In our framework, the activity is assumed to be the same for all working days, which may exhibit slight variations across the working days. A small local maximum at a lag of around 1 average business day (one-fifth of a week in  $\vartheta$ ); a small local maximum at a lag of 2 business days and maxima at 3 and 4 business days also exist. A plausible reason for these remaining autocorrelation peaks is a market-dependent persistence of absolute returns. Autocorrelations with a lag of 1 business day compare with the behaviors of the same market participants, whereas autocorrelations with lags of one half or  $1\frac{1}{2}$  business days compare with the behaviors of different market participants (on opposite sides of the globe). The market-dependent persistence decreases after 2 business days. The predominance of the “meteor shower hypothesis” found by Engle *et al.* (1990) is confirmed by the fact that the autocorrelation curve in



Figure 7.5 does not exhibit strong maxima for each full business day. Yet the remaining small maxima indicate a certain “heat wave” component.

### 7.3.2 Short and Long Memory

The autocorrelation function of volatility decays at a *hyperbolic* rate rather than an exponential rate. In studies based on daily FX prices (e.g., Taylor, 1986) or weekly FX prices (e.g., Diebold, 1988), the number of observations is usually too small for outright rejection of either a hyperbolic or an exponential decay of the autocorrelation functions. In studies with longer daily series such as Ding *et al.* (1993), evidence of long memory is found with the S&P 500 from January 1928 to August 1991 (17,055 observations). To illustrate the presence of the long memory, two curves, one hyperbolic and one exponential, are drawn in Figure 7.5 together with the empirical autocorrelation functions. The hyperbolic curve approximates the autocorrelation function much more closely than the exponential curve. This behavior of volatility is similar to the *fractional noise* process of Mandelbrot and Van Ness (1968) and Mandelbrot (1972), which exhibits hyperbolic decay in the autocorrelation function and thus the long memory serial dependence.

The hyperbolic ( $f_h$ ) and exponential ( $f_e$ ) functions used in the analysis above have the following form:


$$f_h(\tau) = k \tau^{-h}, \text{ and } f_e(\tau) = k e^{-\tau/h} \quad (7.6)$$

where the parameters are  $k$ ,  $h$ , and  $\tau$ .  $\tau$  determines the lag order of the autocorrelation function. The exponential function cannot simultaneously capture the short and long-term persistence, whereas the hyperbolic function is able to capture both successfully. For the hyperbolic function,  $k$  values vary from 0.2 to 0.3 depending on the FX rate, whereas  $h$  is remarkably stable around 0.28 for all the rates.

In Figure 7.4 and the first panel of Figure 7.5, the number of lags are limited to 720 intervals (i.e., 10 days) at the 20-min data frequency. In the second panel of Figure 7.5, the number of lags are extended to 4320 (i.e., 60 days) in  $\vartheta$ -scale. The decay in the volatility autocorrelations is more rapid after 10 days. This type of pattern is not specific to USD-DEM, but is also found in longer time intervals and other FX rates. To explore this behavior further, we compute the autocorrelation function of daily returns (business days) for up to 200 lags and a sample of 20 years. The result is presented in Figure 7.6 and indicates the persistence of the hyperbolic behavior even at the daily frequency.

A process that exhibits a hyperbolic decay in its autocorrelation function is the “fractional noise” of Mandelbrot and Van Ness (1968), which is a purely self-similar fractal. We test the empirical significance for this theoretical process. In Mandelbrot (1972), the autocorrelation function of fractional noise processes is given by

$$a = \frac{|l+1|^{2H} - 2l^{2H} + |l-1|^{2H}}{2} \quad (7.7)$$



**FIGURE 7.6** Autocorrelation function of the absolute business day volatilities in the  $\vartheta$ -time scale. The data are for the USD-DEM rate from June 1, 1973, to June 1, 1993. The hyperbolic (solid curve) and the exponential functions (dotted curve) are superimposed on the empirical autocorrelation function. The 95% confidence intervals are for an identically and independently distributed Gaussian process.

where  $l$  is the lag parameter and  $H$  the Hurst exponent, which lies between 0.5 and 1 for “persistent” fractional noise. For a large number of lags ( $l$ ), the autocorrelation function converges to

$$a \approx H(2H - 1)l^{2(H-1)} \quad (7.8)$$

which has a hyperbolic decay. The autocorrelations of absolute returns in Figures 7.5 and 7.6 also follow a hyperbolic decline. The exponent  $2(H - 1)$  of Equation 7.8 from the USD-DEM volatilities is  $H = 0.87$  in Figure 7.5 and  $H = 0.86$  in Figure 7.6. From the  $H$  values, the factor  $H(2H - 1)$  leads to 0.64 and 0.62, respectively. These values are empirically found to be much lower, which are 0.25 and 0.20, respectively. This indicates that volatility does not follow a *pure* fractional noise process. Volatility is positive definite and has a skewed and fat-tailed distribution, whereas the distribution function of pure fractional noise is Gaussian.

In Peters (1989, 1991), the existence of fractional noise in the returns rather than volatility has been investigated similar to Equation 5.10. These findings claim that a drift exponent different from 0.5 necessarily indicates fractional noise. This conclusion holds only if the distribution forms are stable, but Figure 5.6 does not support this claim. We, therefore, conclude that the return process does not support the fractional noise hypothesis. Unlike volatility, the returns themselves exhibit no significant autocorrelation (see the thin curves in Figures 7.4 and 7.5).

#### 7.4 THE HETEROGENEOUS MARKET HYPOTHESIS

In the earlier sections, we analyzed the presence of two stylized facts. Namely, a hyperbolic decay of the volatility autocorrelations and the “heat wave” effect. Volatility characterizes the market behavior more deeply than just indicating 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 is that, contrary to traditional beliefs, volatility is found to be positively correlated to market presence, activity, and volume. Karpoff (1987), Baillie and Bollerslev (1989), and Müller *et al.* (1990), emphasize the key role of volatility for understanding market structures. The serial correlation studies of LeBaron (1992b,c) show that subsequent returns 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 effect indicates that the forecastability of return is conditional to volatility. Thus, volatility is also an indicator for the persistence of trends.

These 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:

1. Different agents of the heterogeneous market have different time horizons and dealing frequencies. On the side of high dealing frequencies, there 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, there 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.

2. In a homogeneous market, the more agents are present, the faster the price should converge to the “real market value” on which all agents have “rational expectations.” Thus, the volatility should be negatively correlated with market presence and activity. In a heterogeneous market, different market 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 7.3.1, we indicated that the memory in the volatility process is relatively weak at time lags of about  $\frac{1}{2}$  or  $1\frac{1}{2}$  business days when market actors on opposite sides of the globe are related to each other and relatively strong at time lags of about 1 or 2 business days when identical groups of participants are considered.

The market participants of the heterogeneous market hypothesis differ also in other aspects beyond the time horizons and the geographical locations. They may have different degrees of risk aversion, face different institutional constraints, and transaction costs.

#### 7.4.1 Volatilities of Different Time Resolutions

The heterogeneous market hypothesis presented in the previous section is associated with fractal phenomena in the empirical behavior of FX markets. A scaling law relating time horizon and size of price movements (volatility) was identified in Chapter 5. This relation is used here to explain why the perception of volatility differs for market agents with different time horizons.

Short-term traders are constantly watching the market to reevaluate their current positions and execute transactions at a high frequency. Long-term traders may look at the market only once a day or less frequently. A quick price increase of 0.5% followed by a quick decrease of the same size, for example, is a major event for an FX intraday trader but a nonevent for central banks and long-term investors.<sup>4</sup> Long-term traders are interested only in large price movements and these normally happen only over long time intervals (see the scaling law of Müller *et al.*, 1990). Therefore, long-term traders with open positions have no need to watch the market every minute.<sup>5</sup> In other words, they judge the market, its prices, and also its volatility with a coarse time grid. A coarse time grid reflects the view of a long-term trader and a fine time grid that of a short-term trader. Bjorn (1994) follows similar methodologies for building an automatic trading model.

<sup>4</sup> Small, short-term price moves may sometimes have a certain influence on the *timing* of long-term traders' transactions but not on their investment decisions.

<sup>5</sup> They have other means to limit the risk of rare large price movements by stop-loss limits or options.

The time grid in which real traders watch the market is not strictly regular. In the following lagged correlation study, however, we measure volatilities over different but regularly spaced grids. These volatilities are defined in terms of absolute returns. We prefer mean absolute values to roots of mean squares here because they are statistically less dominated by extreme observations, which are rather important in FX markets with their fat-tailed unconditional distribution functions. The convergence of the fourth moment—a requirement for many types of analysis such as the autocorrelation of squared returns—is not guaranteed for empirical returns. In Chapter 5, we demonstrated that the autocorrelations of the returns indicate a stronger signal for powers around one. This argument is reinforced in Dacorogna *et al.* (2001a), where the autocorrelation of absolute returns is also shown to be much more stable under sample size changes than that of the squared returns. Other studies, such as Ding *et al.* (1993), also find absolute returns to be optimal in the autocorrelation studies.

The volatility based on absolute returns has two essential timing parameters (Guillaume *et al.*, 1997):

- The interval size of the time grid in which returns are observed
- The total size of the sample over which it is computed (the number of grid intervals considered)

For exploring the behavior of volatilities of different time resolution, we define two types of volatility. The “coarse” volatility,  $v_c$ , and the “fine” volatility,  $v_f$ , are defined by


$$v_c(t_i) = \left| \sum_{j=1}^n r(\Delta t', t_{i-1} + j \Delta t') \right| \quad \text{and} \quad v_f(t_i) = \sum_{j=1}^n |r(\Delta t', t_{i-1} + j \Delta t')| \quad (7.9)$$

where  $\Delta t' \equiv \Delta t/n$ . Figure 7.7 illustrates this definition where at every time point,  $t_i = t_{i-1} + 6\Delta t'$ , both quantities are simultaneously defined. In this way, the two synchronous time series are obtained whose relation can be explored.

#### 7.4.2 Asymmetric Lead-Lag Correlation of Volatilities

Analyzing the correlation between two time series, such as fine and coarse volatilities, is a standard technique used in empirical finance where the correlation coefficient measures the linear dependence of the two time series. *Lagged correlation* is a more powerful tool to investigate the relation between two time series. The lagged correlation function considers the two series not only simultaneously (at lag 0) but also with a time shift. The correlation coefficient  $\rho_\tau$  of one time series and another one shifted by a positive or negative time lag  $\tau$  is measured and plotted against the value of the lag. The lagged correlation study of this section follows Müller *et al.* (1997a).

Lagged correlation reveals causal relations and information flow structures in the sense of Granger causality. If two time series were generated on the basis of



**FIGURE 7.7** The coarse volatility,  $v_c(t)$ , captures the view and actions of long-term traders while the fine volatility,  $v_f(t)$ , captures the view and actions of short-term traders. The two volatilities are calculated at the same time points and are synchronized.

a synchronous information flow, they would have a symmetric lagged correlation function,  $\rho_\tau = \rho_{-\tau}$ . The symmetry would be violated only by insignificantly small, purely stochastic deviations. As soon as the deviations between  $\rho_\tau$  and  $\rho_{-\tau}$  become significant, there is asymmetry in the information flow and a causal relation that requires an explanation.

In a first analysis, we consider a working-daily time series where weekends are omitted. The variables under study are the “fine volatility” and the “coarse volatility.” Fine volatility is the mean absolute working-daily returns averaged over five observations, so covering a full (working) week. Coarse volatility is the absolute return over a full weekly interval.

The correlation between fine volatility and coarse volatility is a function of the number of lags. When the number of lags is zero, the fine and coarse volatilities are completely identical. In the case of first positive or negative lag, the two intervals do not overlap but follow each other immediately.

The panel on the left hand side of Figure 7.8 shows the lagged correlation function for the USD-DEM in a sample longer than 21 years. The correlation maximum is found at lag zero, which is expected. For the nonzero lags, there is an asymmetry where the coarse volatility predicts fine volatility better than the other way around. The asymmetry is significant for the first two lags where the difference  $\rho_\tau - \rho_{-\tau}$ , represented by the thin curve in Figure 7.8, is distinctly outside the confidence interval for identically and independently distributed observations.

This result can be explained in terms of the heterogeneous market hypothesis presented earlier in this section. For short-term traders, the level of coarse volatility matters because it determines the expected size of trends and thus the scope of trading opportunities. On one hand, short-term traders react to clusters of coarse volatility by changing their trading behavior and so causing clusters of fine volatility. On the other hand, the level of fine volatility does not affect the trading strategies of long-term traders (who often act according to the “fundamentals” of the market).

**TABLE 7.1** Difference between lagged correlation for FX rates and gold.

The sample period is from June 6, 1973, (August 8, 1980 for gold) to February 1, 1995. The lags are measured in weeks and 3 hr (in  $\vartheta$ -time), respectively. The negative values indicate the predictability of finely defined volatility from coarse volatility.

Differences	USD-DEM	USD-JPY	GBP-USD	CHF-USD	DEM-JPY	XAU-USD
Weekly						
$\varrho_1 - \varrho_{-1}$	-0.138	-0.127	-0.130	-0.131	-0.129	-0.122
$\varrho_2 - \varrho_{-2}$	-0.105	-0.047	-0.055	-0.076	-0.074	-0.072
3 hourly						
$\varrho_1 - \varrho_{-1}$	-0.117	-0.136	-0.113	-0.093	-0.100	-0.108
$\varrho_2 - \varrho_{-2}$	-0.058	-0.057	-0.059	-0.056	-0.055	-0.068

Similar behavior of the lagged correlation is observed for other FX rates such as USD-JPY and GBP-USD, cross rates such as DEM-JPY, and gold (XAU-USD). Table 7.1 reports the difference  $\varrho_1 - \varrho_{-1}$  and  $\varrho_2 - \varrho_{-2}$  for a set of these time series. The numbers are similar across the different rates (and also all of the investigated minor FX rates not shown here). The first lag difference is around -0.13 and the second lag difference is around -0.07.


The results with daily data also prevail in high-frequency and in intraday data. Every intra-day study requires an appropriate treatment of the strong intraday seasonality of volatility. Here we use the predefined business time scale  $\vartheta$  presented in Chapter 6. A time series with regular intervals in  $\vartheta$ -time is constructed by selecting the last quote before each point of a regular  $\vartheta$ -grid. As a basic time interval in  $\vartheta$ -time, we choose 30 min. This means there is only some 7 min of physical time during the daily volatility peak in the European afternoon and American morning.<sup>6</sup> Fine volatility is now the mean absolute half-hourly returns averaged over six observations, covering a 3-hr time interval. Coarse volatility is the absolute returns over a full 3-hr interval. All these time intervals are calculated in  $\vartheta$ -time. An interval of 3  $\vartheta$ -hr is clearly smaller than the working day of an FX dealer. It often covers a time span with quite homogeneous market conditions.

Figure 7.8 (right panel) provides the lagged correlation function for USD-DEM in 8 years of half-hour returns. Although the half-hour data cover a shorter time span than the daily series, the number of observations is larger. The findings from the half-hourly data confirm the results from the daily series such that coarse volatility predicts fine volatility. We therefore conclude that these findings are independent of the data frequency.

The intraday behavior of the lagged correlation is similar for other FX rates and gold (see Table 7.1). The empirical findings are similar across the different rates. The first lag difference is around -0.11 and the second lag difference is around -0.06, which are close to the corresponding values of Table 7.1. In the

<sup>6</sup> In fact, a much higher frequency of the series should be avoided due to the fact that price changes observed over 5 min or less can be overly biased by microstructure effects (see Section 7.2).





**FIGURE 7.8** Asymmetric lagged correlation of fine and coarse volatilities for USD-DEM. The *left* figure is for working-daily return in a week. The *right* graph is for high resolution study with half-hourly returns within 3 hr (in  $\vartheta$ -time). The negative lags indicate that the coarse volatility was lagged compared to the fine volatility. The thin curve indicates the asymmetry. The 95% confidence intervals are for identically and independently distributed observations. The sampling period for the left figure is 21 years and 8 months, from June 6, 1973, to February 1, 1995. The sampling period for the right figure is 8 years, from January 1, 1987, to January 1, 1995.

right panel of Figure 7.8, there is also a weak, rather wide local maximum around lag -11, corresponding to -33 hr in  $\vartheta$ -time. This corresponds to a lag of about 1 working day (because a working day is 1/5 rather than 1/7 of a business week). The difference  $q_{\tau} - q_{-\tau}$  also has a significant (negative) peak around lag 11. This effect has been identified in the right panel of Figure 7.8 and discussed in Section 7.3. Following Engle *et al.* (1990), we call it a “heat wave” effect where traders have a better memory of the events approximately 1 working day ago (when they were active) than a broken number working days ago (when other traders on different continents, with different time zones, were active).

The peak around lag -11 can be explained by a residual seasonality that the  $\vartheta$ -scale is unable to capture. However, the  $\vartheta$ -scale is well able to treat *ordinary* seasonality as indicated by the lack of an analogous peak around the positive lag 11. The heat wave effect is more than just seasonality and it cannot be eliminated by a simple time scale transformation. This can be interpreted such that volatility modeling should consider not only volatilities of different time resolutions but also volatilities with the selective memory of individual geographical markets and their time zones.



Assymmetric lead-lag correlation is not only present in the FX market but also in the Eurofutures market as shown in Ballocchi *et al.* (1999a). Figure 7.9 presents the results of a lead/lag correlation analysis for forward rates implied from Euromark contracts on the London International Financial Futures Exchange (LIFFE). The asymmetry is highly significant for the first lag and for all maturities. At lag 1, again coarse volatility predicts fine volatility significantly better than the other way around. The study was conducted with a 3-hr grid in  $\vartheta$ -time where the fine volatility is the mean absolute return measured every 3 hr over 3 days and the coarse volatility is the mean absolute return over the whole 3-day interval. The sample runs from April 1, 1992, to December 30, 1997, which constitutes 700 observations. The effect is rather robust with respect to changes in the definition of the fine and coarse volatilities. Moreover, it is interesting to note that the size of the effect seems to increase when increasing the time-to-start of the forward rate.

To explore this effect on a wider set of parameters, Gilles Zumbach suggested to the following quantities:

$$C(T, n, n') = \text{Corr}(\sigma[T, \frac{T}{2n}](t), \sigma[T, \frac{T}{2n'}](t+T)) \quad (7.10)$$

where  $T = 4$  weeks and  $n$  and  $n'$  are the granularities of our volatility estimator. Then it is possible to compute a quantity  $I$  that depends on both  $n$  and  $n'$ :


$$I(n, n') = C(T, n, n') - C(T, n', n) \quad (7.11)$$

which means that we look at the first lag difference where the lag is 4 weeks. This quantity should in principle be symmetric but we know from Figure 7.8 that it changes sign and is antisymmetric. Figure 7.10 presents the results of a study conducted by Zumbach (private communication), for the quantity  $I$  computed for values of  $n$  going from 2 to 12 over a period  $T$  of 4 weeks on the  $\vartheta$ -time scale. This means that the returns are measured at frequencies as low as 2 weeks to frequencies as high as every 10 min in  $\vartheta$ -time. The FX rate is USD-CHF and the sampling period runs from June 1, 1987, to August 1, 1997. The asymmetry is striking and exists for all these different parameters. The maximum of the effect is obtained for  $n = 11$  for the fine volatility and  $n' = 7$  with differences as high as 0.29 between the two correlations, about two times more than in Table 7.1. Similar figures were also obtained for other FX rates like USD-DEM or USD-JPY.

### 7.4.3 Conditional Predictability


The conditional correlation studies of LeBaron (1992b,c) indicate that subsequent returns are correlated in low-volatility periods and slightly anticorrelated in high-volatility periods. In continuous samples mixed from both low-volatility and high-volatility periods, this important effect indicating the forecastability of return does not exist unconditionally. It exists conditional to volatility. Thus, volatility is also an indicator for the persistence of trends. The idea is to compute the following triplet:

$$(v(t), r(t), r(t + \Delta t))$$



**FIGURE 7.9** Lead-lag correlation of fine and coarse volatilities for four different implied forward rates derived from the Three-Month LIFFE Euromark, with a 3-hr grid in  $\vartheta$ -time. The sampling period is from April 1, 1992, to December 30, 1997. In the panels, a month is represented by  $m$ .

2025 RELEASE UNDER E.O. 14176




**FIGURE 7.10** The correlation difference (Equation 7.11) between coarse and fine volatilities is explored for the USD-CHF FX rate. The asymmetry of the lead-lag correlation (at one lag of 4 weeks) is apparent around the diagonal, which naturally presents a correlation of 1 (and a difference of 0) because we are correlating a quantity with itself. The top half of the graph presents a positive difference in the lagged correlation whereas the bottom half presents the symmetric negative difference. The sampling period is from June 1, 1987, to August 1, 1997. (With permission of Gilles Zumbach.)

where  $v(t)$  is a measure of volatility calculated with the weekly variance of daily returns.<sup>7</sup> Then triplets of similar volatility,  $v(t)$ , are put into the same bin, and the autocorrelation of returns at lag  $\Delta t$ , conditional to  $v(t)$ , is studied:

$$\rho(v) = \rho(r(t), r(t + \Delta t) | v(t)) \quad (7.12)$$

Such an analysis has four parameters,  $\Delta t$  for the returns and the three parameters for the volatility as identified in Section 3.2.4 and in the Equation 3.8:  $\Delta t$ ,  $n$ , and  $p$ .

<sup>7</sup> In principle, any definition of volatility along the lines of Equation 3.8 can be chosen and its parameters varied until the conditional correlation reaches a maximum.



**FIGURE 7.11** The conditional autocorrelation of weekly returns of USD-DEM as a function of the average absolute weekly return over 5 days. The sampling period is June 1, 1973, to June 1, 1994.

This function  $\rho(v)$  is presented for the FX rate USD-DEM on Figure 7.11. It is computed with a  $\Delta t$  of 1 week and a volatility definition that uses the mean weekly absolute returns over 5 weeks. In summary, the parameters for the graph are  $\Delta t = 1$  week,  $n\Delta t = 5$  weeks, and  $p = 2$ . The conditional correlation appears only for data at low frequency. The effect is quite strong for low volatility with a conditional correlation close to 0.3 at its maximum, decreasing down to negative values of -0.15 for high volatility. The computation is done with overlapping bins containing always the same number of observations to avoid changing the significance of the different results. From this figure, it appears that the “current state” of the market changes the price process behavior and the volatility plays an important role beyond its own dynamics. The results shown here for the most traded FX rates are also present in the other FX rates. It was also reported by LeBaron (1992b) for stock indices. Varying the parameters cause this effect to disappear for  $\Delta t$  smaller than 1 day. In the intraday region, influence of the heat wave effect becomes important and overshadows the findings.

# 8

---

## VOLATILITY PROCESSES

### 8.1 INTRODUCTION

One of the many challenges posed by the study of high-frequency data in finance is to build models that can explain the empirical behavior of the data at any frequency at which they are measured from minutes to months. We are now going to examine how conventional models perform when confronted with this problem. In the previous chapter, we discussed the rich structure of the volatility dynamics. We need to introduce new types of volatility models to account for this structure, leading to a higher predictive power.

Many statistical processes proposed in the literature can be described by the following general formula for equally spaced returns  $r_t$ :

$$r_t = \sigma_t \varepsilon_t \quad (8.1)$$

where  $\varepsilon_t$  is an identically and independently distributed (i.i.d.) random variable<sup>1</sup> with zero mean and variance 1. In this chapter,  $t$  denotes the index of a homogeneous time series rather than time itself. “Homogeneous” means equally spaced on any chosen time scale. We usually choose  $\vartheta$ -time as introduced in Chapter 6, so

<sup>1</sup> In this chapter, normal and Student- $t$  distributions of  $\varepsilon_t$  are studied.

the model appropriately accounts for seasonalities. The volatility  $\sigma_t$  is the square root of the variance of the return  $r_t$ .

Many models are based on Equation 8.1, but they largely differ in the modeling of the volatility variable  $\sigma_t$ . We distinguish three main types of volatility modeling:

1. ARCH-type models. These autoregressive conditional heteroskedastic models define the variance  $\sigma_t^2$  of the return  $r_t$  as a function of *past returns*. This function can be simple or rather complicated. In the GARCH process, for example,  $\sigma_t$  also depends on its own past values, but there is always an equivalent formulation that defines  $\sigma_t$  as a function of past returns only. The volatility  $\sigma_t$  is a model variable that cannot be directly observed, but it can be computed if a sufficiently long series of past return values up to  $r_{t-1}$  is known. All the statistical processes discussed in the following sections of this chapter belong to the ARCH type.
2. Stochastic volatility models. In stochastic volatility models, the volatility variable  $\sigma_t$  does not depend on past returns. Instead, it depends on its own past values. The volatility variable  $\sigma_t$  is neither observable nor directly computable from past returns. As a consequence, it is more difficult to estimate the parameters of stochastic volatility models than those of ARCH-type models. The statistical process of  $\sigma_t$  has a memory, so an autoregressive conditional heteroskedastic behavior can be obtained also with stochastic volatility models. There are different types of stochastic volatility models as noted in Taylor (1994); Ghysels and Jasiak (1995), and Ghysels *et al.* (1996). It is possible to model heterogeneous market behavior in the framework of stochastic volatility; a modern example is the cascade model of Ghashghaie *et al.* (1996) and Breyman *et al.* (2000) where volatility modeling is inspired by turbulence models.
3. Models based on realized volatility. Rather than modeling  $\sigma_t$ , (Andersen *et al.*, 2000) propose to define  $\sigma_t$  as the realized volatility computed at index  $t - 1$ . This realized volatility is computed with high-frequency data, with return intervals of, for example, 30-min, in order to keep stochastic errors low. The time interval of the main model (i.e., the interval between the indices  $t - 1$  and  $t$ ) is usually much larger (e.g., 1 working day). This means using realized volatility at  $t - 1$  as a predictor of the volatility between  $t - 1$  and  $t$  by relying on the volatility clustering. This model has the advantage of using empirical data instead of model assumptions that might be wrong. However, it has some disadvantages:
  - Realized volatility is biased if computed at high frequency (see Section 7.2). A bias correction method such as Equation 7.4 would improve the model.
  - Realized volatility computed at high frequency (fine volatility) lags behind coarse volatility in the lead-lag analysis (see Section 7.4.2). This lag leads to suboptimal forecast quality when predicting the volatility of the next step of the model.

- In general, realized volatility at  $t - 1$  may not be the best predictor of volatility between  $t - 1$  and  $t$ . It should be replaced by a more sophisticated *forecast* of realized volatility at  $t$ .

In most of these statistical processes of  $\sigma_t^2$ , it is possible to add some terms modeling external (exogenous) influences. If volume figures at  $t - 1$  are available, for example, they may be a piece of information to predict the volatility  $\sigma_t$ . The processes discussed here are not of this type, they are univariate.

In the remainder of this chapter, we stay within the framework of ARCH-type modeling and compare different models. The ultimate quality criterion of a model is its predictive power. Therefore there are some volatility forecast tests in Section 8.4. Forecasting is further discussed in Chapter 9 where it is the main subject.

## 8.2 INTRADAY VOLATILITY AND GARCH MODELS

The Autoregressive Conditional Heteroskedastic (ARCH) model of Engle (1982) and its generalized version (GARCH) by Bollerslev (1986) are widely used, not only in the foreign exchange (FX) literature (see, for a review, Bollerslev *et al.*, 1992) but also as the basic framework for empirical studies of the market microstructure such as the impact of news (Goodhart and Figliuoli, 1991; Goodhart *et al.*, 1993) and central bank interventions (Goodhart and Hesse, 1993; Peiers, 1997), or inter and intramarket relationships in Engle *et al.* (1990) and Baillie and Bollerslev (1990). The main assumption behind this class of models is the relative homogeneity of the price discovery process among market participants at the origin of the volatility process. In other words, the conditional density of one GARCH process is assumed to adequately capture all the information and the news. In particular, GARCH parameters for the weekly frequency theoretically derived from daily empirical estimates are usually within the confidence interval of weekly empirical estimates (Drost and Nijman, 1993).

However, we have already seen in this book that several empirical facts are at odds with this homogeneous view of the market. First, the long memory of the volatility (Section 7.3.2) indicates the presence of several market components corresponding to several time horizons. Note that this property of the volatility has already been successfully incorporated in the GARCH setting as the fractionally integrated GARCH (Baillie *et al.*, 1996). Second, at the intradaily frequency, round-the-clock time series reveal seasonal patterns that reflect, among others, the geographical dispersion of the traders, concentrated in three main geographical areas: Asia, Europe, and America. Although the first investigations of the effect of these different geographical locations seemed to indicate that news would uniformly spread out around the world (the so-called meteor shower hypothesis in Engle *et al.*, 1990), we saw traces of heat wave effects in the previous chapter. Third, exchange rates movements are not necessarily related to the arrival of news when inspected at the intraday frequency, Goodhart (1989), reflecting the fact that intraday traders may have other constraints and objectives than, for

example, longer-term traders. Fourth, at extremely high frequencies, FX rates exhibit distinct microstructure effects due to the price formation process as studied in Chapter 5.

In this section, we investigate the importance of this heterogeneity for the modeling of the foreign exchange (FX) markets using the GARCH setting. More specifically, we show that estimates of a GARCH process with data in physical time are likely to be spurious, even though estimates for one particular frequency seem to be reasonable. Estimates are only consistent when the seasonal patterns are taken into account. However, even when these seasonal patterns are accounted for, the aggregation properties of the GARCH model break down at the intraday frequencies, revealing the presence of traders with different risk profiles. In addition to the presence of different trader categories, we observe microstructure effects when analyzing returns over time intervals shorter than about 90 min. At the other extreme, the instability of coefficient estimates over different subperiods of 6 months suggests the presence of seemingly random long-term fluctuations. Finally, these misspecifications of the GARCH process result in its quite poor out-of-sample predictive power for the volatility as compared to realized volatility.

### 8.2.1 Parameter Estimation of GARCH Models

The GARCH(1,1) process is defined as follows:

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \quad (8.2)$$

where  $\sigma_t^2$  is the conditional variance and  $\varepsilon_t^2$  is the squared innovation.

To test the effects of the temporal heterogeneity of the markets, this GARCH(1,1) process is estimated for several frequencies. The lowest analyzed frequency is daily and the highest frequency is defined by a homogeneous time series with 10-min intervals. At the higher frequencies (intervals less than 2 hr), we include a fourth-order autoregressive (AR(4)) term  $\mu_t = \sum_{i=1}^4 \phi_i r_{t-i}$  in Equation 8.1 to account for the statistically significant (negative) autocorrelation of the returns at these frequencies (see Section 5.2.1). The regression equation for the return process is

$$r_t = \mu_t + \varepsilon_t \quad (8.3)$$

At lower frequencies such a term is not needed, and we use the process of Equation 8.1.

The parameters of the process are estimated as follows. Let  $\theta$  denote the set of parameters characterizing the process. Assuming that the innovations  $\varepsilon_t$  are normally distributed, the log-likelihood function is

$$\mathcal{L}(\theta) = -\frac{n}{2} \ln(2\pi) - \frac{1}{2} \sum_{i=1}^n \left( \ln(\sigma_i^2) + \frac{\varepsilon_i^2}{\sigma_i^2} \right) \quad (8.4)$$

where the index  $t$  has been substituted by  $i$ . The number of observations used for the estimation is  $n$ . An initial fraction of data must be reserved and used for the

Downloaded from https://www.cambridge.org/core. University of Cambridge, on 05 Jun 2020 at 10:00:00, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9780511527813.010



build-up of  $\sigma_t^2$ , because of the memory of the volatility process. An estimate  $\tilde{\theta}$  for the parameters is given by the solution of the maximization problem

$$\max_{\theta} \mathcal{L}(\theta)$$

The log-likelihood procedure has many desirable properties.<sup>2</sup> The solution is independent of the coordinate system in which the parameters are defined, such that the estimation can be done in any parametrization and the results will be identical, up to the chosen parameter transformation. This property is true for finite samples and any data set, assuming a non-degenerate maximum. Even if the process is misspecified (i.e., the data were not generated by the estimated process), the maximum is identical in any coordinate system. Estimating GARCH processes by maximum likelihood is difficult because of the presence of a one-dimensional manifold in the parameter space where the likelihood function is large and almost constant (for a discussion of this point and a good practical solution using the property mentioned above, see Zumbach, 2000).

The assumption of conditional normality can be relaxed by assuming a Student- $t$  distribution for  $\varepsilon_t$  (Baillie and Bollerslev, 1989) or the generalized exponential distribution (Nelson, 1991). Both of these distributions have fat tails. In the case of the Student- $t$  distribution, the log-likelihood function takes the following form:

$$\begin{aligned} \mathcal{L}(\theta) = & -\frac{n}{2} \left[ \ln(\nu - 2) + 2 \ln \left[ \pi^{1/2} \Gamma \left( \frac{\nu}{2} \right) \right] - \ln \Gamma \left( \frac{\nu+1}{2} \right) \right] \\ & - \frac{1}{2} \sum_{i=1}^n \left[ \ln(\sigma_i^2) + (\nu + 1) \ln \left( 1 + \frac{\varepsilon_i^2}{\sigma_i^2(\nu-2)} \right) \right] \end{aligned} \quad (8.5)$$

where  $\nu$  is the number of degrees of freedom of the Student- $t$  distribution and  $\Gamma$  is the usual gamma function. Both forms of the log-likelihood function are valid for any process following Equation 8.1, not only GARCH but also the process we shall study in Section 8.3.1.

The maximum of the likelihood function is found by an iterative procedure that combines two methods: a genetic algorithm (GA) (Goldberg, 1989; Pictet *et al.*, 1995) and the Berndt, Hall, Hall, and Hausman (BHHH) algorithm (Berndt *et al.*, 1974) which is a variant of the gradient descent method. The initial solutions are chosen randomly to avoid any a priori bias in the estimation and stored in “genes,” which form an initial population. Starting from this population, the genetic algorithm constructs a new population using its selection and reproduction method (Pictet *et al.*, 1995). The solutions with the highest log-likelihood found by the genetic algorithm are used as starting points of the BHHH algorithm, which leads to a further improvement. Once convergence of the BHHH is achieved, the next generation of the GA is computed on the basis of the previous solutions

<sup>2</sup> See Davidson and MacKinnon (1993) for a general reference.

obtained with the BHHH algorithm and a set of solutions from the previous generation. This iterative procedure continues until no improvement of the solution is found. The BHHH algorithm alone can be trapped in local maxima of the log-likelihood instead of finding the global maximum. The chosen combination with a genetic algorithm has the advantage of avoiding local maxima. The method is rather fast, notwithstanding the very large number of observations (368,000 data points for the 10-min frequency). Robust standard errors are computed using the variance-covariance matrix estimation of White (1980).

### 8.2.2 Temporal Aggregation of GARCH Models

If the empirical data can be described as generated by one GARCH(1,1) process at one particular data frequency, the behavior of the data sampled at any other frequency is theoretically determined by temporal aggregation (or disaggregation) of the original process. These theoretically derived processes at different frequencies can be compared to the empirically estimated processes at the same frequencies. Significant deviations between empirical and theoretical results lead to the rejection of the hypothesis of only one GARCH process. We can show then that there is more than one relevant frequency in the volatility generation, and the market can be called temporally heterogeneous, as already found in Section 7.4.

There are two approaches for the theoretical aggregation of GARCH models. The GARCH model can be viewed as either a jump process (Drost and Nijman, 1993) or a diffusion process (Nelson and Foster, 1994). Both approaches lead to very similar results, so we only report results based on Drost and Nijman (1993). In both approaches, the sum of  $\alpha_1$  and  $\beta_1$  (of Equation 8.2) tends to 1 as the frequency increases. The autoregressive parameter  $\beta_1$  tends to 1, whereas the moving average parameter  $\alpha_1$  tends to 0. In other words, the higher the frequency, the longer the clusters of volatility as measured in numbers of time series intervals.

Because previous results confirmed the adequacy of these theoretical results at the daily and weekly frequencies (Drost and Nijman, 1993), we use the daily estimations as a starting point to compute the results for the higher frequencies. High frequencies also have the advantage of high statistical significance.

Drost and Nijman (1993) show that symmetric weak GARCH(1,1) processes are closed under temporal aggregation. A process is symmetric if the marginal distribution of returns is symmetric. The term “weak GARCH(1,1)” is exactly defined by Drost and Nijman (1993). It encompasses all processes that essentially follow Equation 8.2 with some weak, nonlinear deviations that are not visible in the autocorrelation of volatility. More precisely, if  $\varepsilon_t$  is a symmetric weak GARCH(1,1), following the equation  $\sigma_t^2 = \alpha_0 + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$ , then the high-frequency parameters  $\alpha_0$ ,  $\alpha$ , and  $\beta$  and the kurtosis  $\kappa_\varepsilon = E[\varepsilon_t^4]/(E[\varepsilon_t^2])^2$  determine the corresponding low-frequency parameters. We obtain the symmetric weak GARCH(1,1) process  $\bar{\varepsilon}_{(m)tm}$ , with

$$\overline{\sigma^2}_{(m)tm} = \bar{\alpha}_0(m) + \bar{\alpha}_{(m)}\bar{\varepsilon}_{(m)tm-m}^2 + \bar{\beta}_{(m)}\overline{\sigma^2}_{(m)tm-m} \quad (8.6)$$

and kurtosis  $\bar{\kappa}_{(m)\varepsilon}$  where

$$\bar{\alpha}_{0(m)} = m \alpha_0 \frac{1 - (\beta + \alpha)^m}{1 - (\beta + \alpha)} \quad (8.7)$$

$$\bar{\alpha}_{(m)} = (\beta + \alpha)^m - \bar{\beta}_{(m)} \quad (8.8)$$

$$\begin{aligned} \bar{\kappa}_{(m)\varepsilon} = & \quad 3 + (\kappa_\varepsilon - 3)/m + 6(\kappa_\varepsilon - 1) \\ & \times \frac{\{m - 1 - m(\beta + \alpha) + (\beta + \alpha)^m\} \{\alpha - \beta\alpha(\beta + \alpha)\}}{m^2(1 - \beta - \alpha)^2(1 - \beta^2 - 2\beta\alpha)} \end{aligned} \quad (8.9)$$

$|\bar{\beta}_{(m)}| < 1$  is the solution of the quadratic equation

$$\frac{\bar{\beta}_{(m)}}{1 + \bar{\beta}_{(m)}^2} = \frac{a(\beta, \alpha, \kappa_\varepsilon, m)(\beta + \alpha)^m - b(\beta, \alpha, m)}{a(\beta, \alpha, \kappa_\varepsilon, m)\{1 + (\beta + \alpha)^{2m}\} - 2b(\beta, \alpha, m)} \quad (8.10)$$

with

$$a(\beta, \alpha, \kappa_\varepsilon, m) = \quad (8.11)$$


$$\begin{aligned} & m(1 - \beta)^2 + 2m(m - 1) \frac{(1 - \beta - \alpha)^2(1 - \beta^2 - 2\beta\alpha)}{(\kappa_\varepsilon - 1)\{1 - (\beta + \alpha)^2\}} \\ & + 4 \frac{\{m - 1 - m(\beta + \alpha) + (\beta + \alpha)^m\} \{\alpha - \beta\alpha(\beta + \alpha)\}}{1 - (\beta + \alpha)^2} \end{aligned}$$

$$b(\beta, \alpha, m) = \quad \{\alpha - \beta\alpha(\beta + \alpha)\} \frac{1 - (\beta + \alpha)^{2m}}{1 - (\beta + \alpha)^2} \quad (8.12)$$

These formulas are used to determine the parameters of the aggregated GARCH processes and can also be used for going from low to high-frequency (i.e., for disaggregation).

When exploring temporal aggregation, we have to choose a time scale. Seasonality is not the subject of an aggregation study, but might disturb it. Eliminating seasonalities by using the  $\vartheta$ -scale presented in Chapter 6 is a natural choice. However, we have additionally tested an alternative time scale which we call a business time scale in the remainder of this section. This business time simply omits the weekend periods from Friday 22:30 GMT to Sunday 22:30 GMT, when markets are virtually closed.

As a third time scale, we have tried physical time. In physical time, weekends cover two-sevenths of the total sample. This causes a complete breakdown of the estimation procedure, yielding very large  $\alpha_1$  estimates. Physical time including weekends is simply unusable here. The aforementioned business time is a usable substitute of physical time from which it only differs in its omission of weekends.




**FIGURE 8.1** Aggregation of the GARCH(1,1) for estimated coefficients in business time (●) and theoretically derived coefficients (Δ) using the (Drost and Nijman, 1993) results for USD-DEM, for different aggregation factors (1 = 10 min; 2 = 20 min; 3 = 30 min; 6 = 1 hr; 12 = 2 hr; 36 = 6 hr; 72 = 12 hr; 144 = 24 hr). The labels of the estimated coefficients (●) are printed in bold. The diagonal dotted line represents the stationarity limit for which  $\alpha_1 + \beta_1 = 1$ . Sampling period: 7 years from January 1, 1987, to December 31, 1993.

### 8.2.3 Estimates of GARCH(1,1) for Various Frequencies

Time series of USD-DEM have been sampled with frequencies between 10 min and 1 day. For each series, the GARCH(1,1) coefficients have been estimated using the procedure of Section 8.2.1. The resulting coefficients  $\alpha_1$  and  $\beta_1$  (see Equation 8.2) are plotted in Figure 8.1 in the form of black circles, which are labeled by the number of 10-min intervals contained in the time series interval. The label “144” thus means daily sampling.

For comparison, the theoretical values of  $\alpha_1$  and  $\beta_1$  are also plotted as triangles. The reference values at daily frequency (label 144) are estimated from real data, but the values at all other frequencies are computed from these reference values according to Drost and Nijman (1993), as explained in Section 8.2.2. A computation according to Nelson and Foster (1994) yields similar results that are not plotted here.



**FIGURE 8.2** Aggregation of the GARCH(1,1) for estimated coefficients in  $\vartheta$ -time (●) and theoretically derived coefficients (△) using the (Drost and Nijman, 1993) results for the USD-DEM for different aggregation factors (1 = 10 min; 2 = 20 min; 3 = 30 min; 6 = 1 hr; 12 = 2 hr; 36 = 6 hr; 72 = 12 hr; 144 = 24 hr). The labels of the estimated coefficients (●) are printed in bold. The diagonal dashed line represents the limit for which  $\alpha_1 + \beta_1 = 1$ . Sampling period: 7 years from January 1, 1987, to December 31, 1993.

Although the coefficient estimates may look quite reasonable for some lower frequencies, the global picture for all frequencies appears quite odd. In particular, the  $\beta_1$  estimates for frequencies higher than 2 hr decrease down to values close to 0.75, whereas the theory, represented by the triangles in Figure 8.1, suggests that  $\beta_1$  should tend to one. The triangles are very far from the corresponding black circles. The hypothesis of volatility being generated by only one GARCH(1,1) process is clearly rejected with the high significance of high-frequency data analyzed over 7 years.

The results of Figure 8.1 are computed on the basis of the business time introduced at the end of Section 8.2.2. Figure 8.2 shows the corresponding results based on  $\vartheta$ -time. The time scale  $\vartheta$  is in fact a better choice because of its better deseasonalization. However, the results of Figure 8.2 are similar to those of Figure 8.1. The strong deviation between theoretical and empirically estimated

"FOOD" IN THE "FOOD"

**TABLE 8.1** Results of the GARCH(1,1) estimation in business time.

GARCH(1,1) parameter estimates for USD-DEM, using the business time scale, for different frequencies. Robust standard errors are given in parentheses. The coefficients with a prime are computed from the (dis)aggregation formulas for the jump hypothesis of Drost and Nijman (1993). The daily interval serves as a reference basis.

Interval	$\alpha_0$	$\alpha_1$	$\beta_1$	$\alpha_1 + \beta_1$	$\alpha'_1$	$\beta'_1$	$\alpha'_1 + \beta'_1$
10 min	$2.15 \cdot 10^{-8}$ ( $0.17 \cdot 10^{-8}$ )	0.227 (0.0013)	0.752 (0.0012)	0.979	0.001	0.999	1.000
20 min	$2.66 \cdot 10^{-8}$ ( $0.15 \cdot 10^{-8}$ )	0.179 (0.0037)	0.816 (0.0051)	0.995	0.002	0.997	0.999
30 min	$2.65 \cdot 10^{-8}$ ( $0.18 \cdot 10^{-8}$ )	0.143 (0.0062)	0.853 (0.0101)	0.996	0.003	0.996	0.999
1 hr	$1.79 \cdot 10^{-7}$ ( $0.42 \cdot 10^{-7}$ )	0.142 (0.0066)	0.784 (0.0114)	0.926	0.006	0.992	0.999
2 hr	$3.11 \cdot 10^{-8}$ ( $0.13 \cdot 10^{-8}$ )	0.023 (0.0020)	0.970 (0.0015)	0.993	0.011	0.986	0.997
6 hr	$2.43 \cdot 10^{-7}$ ( $0.22 \cdot 10^{-7}$ )	0.041 (0.0039)	0.941 (0.0061)	0.982	0.029	0.962	0.991
12 hr	$1.07 \cdot 10^{-6}$ ( $0.17 \cdot 10^{-6}$ )	0.054 (0.0061)	0.905 (0.0102)	0.959	0.046	0.936	0.982
24 hr	$1.91 \cdot 10^{-6}$ ( $0.67 \cdot 10^{-6}$ )	0.068 (0.0095)	0.897 (0.0153)	0.965			

coefficients already starts with the 6-hr frequency. The conclusions on temporal aggregation of GARCH are the same. The choice of the time scale has no strong impact on a temporal aggregation study, as long as physical time with its high weight of weekends is avoided.

Detailed results for the two time scales are also listed in Tables 8.1 and 8.2. Table 8.1 presents the results obtained for USD-DEM in the business time scale and Table 8.2 for the same rate but in the  $\vartheta$ -time scale. The error estimates of the results provide more evidence against the hypothesis of only one GARCH(1,1) process generating the data. Even the theoretically computed coefficients at low frequency, which seem quite close to the estimated coefficients, are often outside the confidence intervals. Only the coefficients for the GARCH process are provided in Tables 8.1 and 8.2, even when an AR(4) term was included in Equation 8.1 for frequencies higher than 2 hr (as discussed at the beginning of Section 8.2.1). We have observed that the inclusion of this autoregressive term in the return equation does not significantly change the values of the GARCH coefficients.

The coefficient estimates are quite similar across different FX rates.<sup>3</sup> The hypothesis of only one GARCH(1,1) process is rejected for all the FX rates we tested, not only USD-DEM. The volatility clusters have about the same size—if measured in numbers of time series intervals—for all levels of aggregation. In

<sup>3</sup> See Andersen and Bollerslev (1997b) for similar results.

**TABLE 8.2** Results of the GARCH(1,1) estimation in  $\vartheta$ -time.

GARCH(1,1) parameter estimates for USD-DEM, using  $\vartheta$ -time, for different frequencies. Robust standard errors are given in parentheses. The coefficients with a prime are computed from the (dis)aggregation formulas for the jump hypothesis of Drost and Nijman (1993). The daily interval serves as a reference basis.


Interval	$\alpha_0$	$\alpha_1$	$\beta_1$	$\alpha_1 + \beta_1$	$\alpha'_1$	$\beta'_1$	$\alpha'_1 + \beta'_1$
10 min	$4.09 \cdot 10^{-9}$ ( $0.27 \cdot 10^{-9}$ )	0.153 (0.0047)	0.839 (0.0049)	0.992	0.001	0.999	1.000
20 min	$1.24 \cdot 10^{-8}$ ( $0.84 \cdot 10^{-8}$ )	0.149 (0.0057)	0.830 (0.0063)	0.979	0.001	0.998	0.999
30 min	$2.56 \cdot 10^{-8}$ ( $0.21 \cdot 10^{-8}$ )	0.153 (0.0077)	0.815 (0.0091)	0.968	0.002	0.997	0.999
1 hr	$1.36 \cdot 10^{-8}$ ( $0.46 \cdot 10^{-8}$ )	0.047 (0.0094)	0.942 (0.0129)	0.988	0.004	0.995	0.999
2 hr	$1.65 \cdot 10^{-8}$ ( $0.28 \cdot 10^{-8}$ )	0.031 (0.0014)	0.962 (0.0022)	0.993	0.008	0.989	0.997
6 hr	$5.93 \cdot 10^{-8}$ ( $0.40 \cdot 10^{-8}$ )	0.029 (0.0011)	0.963 (0.0013)	0.992	0.023	0.971	0.994
12 hr	$1.91 \cdot 10^{-7}$ ( $0.45 \cdot 10^{-7}$ )	0.039 (0.0013)	0.948 (0.0047)	0.987	0.038	0.949	0.988
24 hr	$8.08 \cdot 10^{-7}$ ( $2.74 \cdot 10^{-7}$ )	0.061 (0.0119)	0.915 (0.0155)	0.975			

other words, the volatility memory seems quite short-lived when measured with high-frequency data and long-lived when measured with data of daily or lower frequency. The information content of the volatility variable  $\sigma_t$  is not the same for different frequencies. Different volatilities are relevant at different frequencies. We attribute this, along with other authors (Andersen and Bollerslev, 1997a), to the presence of many independent volatility components in the data. This is again the signature of market heterogeneity. The GARCH model does not capture the heterogeneity of traders acting at different time horizons.

This is a plausible explanation of the abnormal results we obtain at high frequencies from the estimation of the GARCH model using a Student- $t$  distribution instead of the normal distribution such as in Baillie and Bollerslev (1989).<sup>4</sup> GARCH is misspecified, no matter which form of the conditional distribution of returns is chosen.

To further assess the behavior of the volatility as estimated by GARCH(1,1) processes, we have investigated the temporal stability of the coefficient estimates for several subsamples. Figure 8.3 provides the estimations of the GARCH parameters for USD-DEM at the 2-hr time interval, using  $\vartheta$ -time, for subsamples of 6 months, with about 2,190 observations per subsample. As can be seen, the

<sup>4</sup> Although the algorithm converges, the sum of the  $\alpha_1$  and  $\beta_1$  increasingly exceeds 1 as the frequency becomes higher. One also finds excess residual skewness and kurtosis. Since these results are robust to the size of the sample, they cannot be attributed to a larger number of tail observations.



**FIGURE 8.3** Temporal stability of the GARCH(1,1) coefficients for subperiods of 6 months for the USD-DEM at the 2-hr frequency. The time scale is  $\hat{v}$ -time. In the parameter space, the coefficients are represented by black circles ( $\bullet$ ) and connected by lines indicating the temporal sequence. Sampling period: 7 years from January 1, 1987, to December 31, 1993.

coefficients are not stable over time. Some of them are in the left half of Figure 8.3, quite far from the others. Moreover, these aberrant coefficients are not directly connected in the temporal sequence. The shifts in coefficient values have an irregular sequence in time, as shown by the lines connecting the points in Figure 8.3. The hypothesis of all parameters being equal across the subsamples can be rejected by using a likelihood-ratio test (see e.g., Hamilton, 1994) with very high significance. This is again a sign of misspecification of the model, but it may also indicate changes in market behavior.

The forecasting quality of GARCH models will be tested in Section 8.4.2. There we shall see that an increasing sample size does not improve the volatility forecasts from GARCH models. The forecasting quality saturates when increasing the sample size after a certain threshold value. The subsamples used for Figure 8.3 are large enough, so the erratic behavior of GARCH coefficients in that figure cannot be attributed to small sample sizes.



### 8.3 MODELING HETEROGENEOUS VOLATILITIES

In Chapter 7 we showed that there is asymmetry in the interaction between volatilities measured at different frequencies. A coarsely defined volatility predicts a fine volatility better than the other way around. This effect is not present in a simple GARCH model. More complex types of ARCH models have to be developed to account for the heterogeneity found in high-frequency data, such as the HARCH (Heterogeneous Autoregressive Conditional Heteroskedasticity) model.

The HARCH process proposed in this section has a variance equation based on returns over intervals of *different sizes*. The empirical behavior of lagged correlation can be reproduced well with this new process. At the same time, HARCH is able to reproduce the long memory of volatility,<sup>5</sup> as found in Section 7.3.2, Dacorogna *et al.* (1993), and Ding *et al.* (1993). Moreover, the terms of the conditional variance of HARCH reflect the component structure of the market in a natural way.

As with most processes from the ARCH family, HARCH is based on squared returns,<sup>6</sup> with their good analytical tractability. Whereas the convergence problem of the fourth moment of the return distribution forced us to define volatility in terms of absolute returns in the correlation analysis of Section 7.4.1, there is no such constraint for the volatility equation of the HARCH process.

#### 8.3.1 The HARCH Model

In this section, we present the HARCH model as it was first presented by Müller *et al.* (1997a). This should facilitate the understanding of the approach but this initial formulation as the initial ARCH formulation is cumbersome to compute. In the next sections, we shall see a formulation with a much faster and simpler computation and estimation. As in Equation 8.2, the returns  $r_t$  of a HARCH( $n$ ) process are defined with the random variable  $\varepsilon_t$ , which is identically and independently distributed (i.i.d.) and follows a distribution function with zero expectation and unit variance:<sup>7</sup>

$$r_t = \sigma_t \varepsilon_t$$

$$\sigma_t^2 = c_0 + \sum_{j=1}^n c_j \left( \sum_{i=1}^j r_{t-i} \right)^2 \quad (8.13)$$

where

$$c_0 > 0, \quad c_n > 0, \quad c_j \geq 0 \text{ for } j = 1 \dots n-1 \quad (8.14)$$

<sup>5</sup> The FIGARCH process, Baillie *et al.* (1996), has been designed to model the long memory but cannot reproduce the lead-lag correlation of Section 7.4.2 as it is still based on returns observed over intervals of constant size.

<sup>6</sup> Except for a process class proposed by Ding *et al.* (1993), which models volatility in terms of different powers of absolute returns

<sup>7</sup> Here, we utilize a normal distribution and alternatively explore Student- $t$  distributions.

The equation for the variance  $\sigma_t^2$  is a linear combination of the squares of *aggregated* returns. Aggregated returns may extend over some long intervals from a time point in the distant past up to time  $t - 1$ . The heterogeneous set of relevant interval sizes leads to the process named HARCH for “Heterogeneous Autoregressive Conditional Heteroskedasticity.” The first “H” may also stand for the heterogeneous market if we follow that hypothesis as proposed in Section 7.4.

The HARCH process belongs to the wide ARCH family but differs from all other ARCH-type processes in the unique property of considering the volatilities of returns measured over different interval sizes. The Quadratic ARCH (QARCH) process (see Sentana, 1991) is an exception. Although QARCH was not developed for treating different interval sizes, it can be regarded as a generalized form of HARCH as explained in Section 8.3.3.

The coefficients  $c_1 \dots c_n$  should not be regarded as free parameters of the model. The heterogeneous market approach leads to a low number of free model parameters, which determine a much higher number  $n$  of coefficients modeling the long memory of volatility.

The explicit formulation of HARCH(2) may help to illustrate the special properties of the HARCH process.<sup>8</sup> The variance equation of HARCH(2) can be written in two forms:

$$\begin{aligned} \sigma_t^2 &= c_0 + c_1 r_{t-1}^2 + c_2 (r_{t-1} + r_{t-2})^2 = \\ &c_0 + (c_1 + c_2) r_{t-1}^2 + c_2 r_{t-2}^2 + 2 c_2 r_{t-1} r_{t-2} \end{aligned} \quad (8.15)$$

The second form of this HARCH(2) process can be identified as an ordinary ARCH(2) process, except for its last term which contains the mixed product  $r_{t-1}r_{t-2}$ . In other ARCH-type processes, the absolute values of returns matter where in HARCH, also their *signs* matter. Two subsequent returns of the same size and in the same direction will cause a higher contribution to the variance process than two subsequent returns that cancel out each other.

The variance, the unconditional expectation of squared returns, can be derived from Equation 8.13:

$$E(r_t^2) = E(\sigma_t^2) = c_0 + \sum_{j=1}^n c_j \left[ \sum_{i=1}^j E(r_{t-i}^2) \right] \quad (8.16)$$

The cross products, such as  $r_{t-1}r_{t-2}$  in Equation 8.15, have no influence here as their expectation is zero. A necessary condition of stationarity is constant unconditional variance:

$$E(r_t^2) = E(r_{t-i}^2), \quad i \geq 1 \quad (8.17)$$

<sup>8</sup> Whereas HARCH(1) is identical to ARCH(1).

Inserting this in Equation 8.16, we obtain the variance

$$E(r_t^2) = \frac{c_0}{1 - \sum_{j=1}^n j c_j} \quad (8.18)$$

which must be finite and positive

$$\sum_{j=1}^n j c_j < 1 \quad (8.19)$$

This necessary stationarity condition is also a sufficient condition for both the stationarity of the process and the existence of the variance, the second moment. Proving this is not trivial and we do not follow here the path chosen by Engle (1982) and Bollerslev (1986) because the mixed products, such as  $r_{t-1}r_{t-2}$ , make the matrix formulation of the problem difficult. The HARCH process can be seen as a Markov chain. Meyn and Tweedie (1993) have obtained some results for the ergodicity and recurrence of Markov chains that can be used for proving the stationarity and the moment condition. The complete proof is given by Dacorogna *et al.* (1996).


The conditions for the existence and constant unconditional expectation of higher moments can be obtained through steps analogous to Equations 8.16 through 8.19, but the computation becomes increasingly tedious for higher moments and larger  $n$  values. The expectation of the  $2m^{\text{th}}$  moment is

$$E(r_t^{2m}) = E(\sigma_t^{2m}) E(\varepsilon_t^{2m}) \quad (8.20)$$

Equation 8.13 of the variance is inserted in  $E(\sigma_t^{2m})$  and all the terms are explicitly computed. Some products of powers of returns have nonzero expectations, leading to an equation system for these expectations and  $E(\sigma_t^{2m})$ . The equation system has the dimension  $m$  for  $n = 2$  and higher for larger  $n$  values. In the relatively simple case of the fourth moment ( $m = 2$ ) of HARCH(2) ( $n = 2$ ), the expectation  $E(r_t^2 r_{t-1}^2)$  has to be computed and solved together with the equation for  $E(r_t^4)$ . In the standard case of  $\varepsilon_t$  following a normal distribution,  $N(0, 1)$ ,  $E(\varepsilon_t^4) = 3$  and the following necessary condition is obtained to keep the fourth moment finite

$$3 [c_2^2 + (c_1 + c_2)^2] + c_2 [1 + 3 (c_1^2 + 6c_1c_2 + 4c_2^2)] < 1 \quad (8.21)$$

The sufficiency of this necessary fourth moment condition is also proven in Dacorogna *et al.* (1996). In Figure 8.4, the second and fourth moment conditions according to Equations 8.19 and 8.21, plus the sixth moment condition following an analogously derived equation, are plotted for a HARCH(2) process. Processes with a finite second and a diverging fourth moment exist in a large part of the  $c_1$ - $c_2$ -plane. This is remarkable because half-hourly FX returns have an empirical distribution with a tail index between 2 and 4, as found in Section 5.4.2 and Dacorogna *et al.* (2001a).



**FIGURE 8.4** Moment conditions for the coefficients  $c_1$  and  $c_2$  of a HARCH(2) process with a normally distributed  $\varepsilon_t$ . The straight line on the right represents the boundary for the stationarity and the existence of the second moment. The curves in the middle and on the left represent the boundaries for the existence of the 4th and 6th moments.

**8.3.2 HARCH and Market Components**

The memory in the volatility is known to be long, as discussed in Section 7.3.2. Therefore, we need a high order of HARCH, a large  $n$ , to model the behavior of empirical time series. This implies a high number of coefficients  $c_j$ , which should not be free parameters of HARCH. We need a parsimonious parametrization. In the case of ARCH, some high-order processes can be formulated as low-order GARCH processes (Bollerslev, 1986), but no analogous method is at hand to reduce the number of HARCH parameters.

Our approach of parsimonious parametrization allows for the exploration of the component structure of the market. The coefficients  $c_j$  reflect the relative impact of different market components with different relevant time intervals. Therefore, we select  $m$  market components corresponding to  $m$  free parameters, each associated to some coefficients  $c_j$  in a limited range of  $j$ . The  $j$  ranges are separated by powers of a natural number  $p$ , so the typical time interval size of a component differs from that of the neighbor component by a factor of  $p$ . All  $c_j$  values within one component are assumed to be the same:

$$c_j = C_i = C_{i(j)} \\ i(j) = \max \left( k \mid k \in \mathbb{N} \wedge k < \frac{\log j}{\log p} + 2 \right), \quad j = 1 \dots p^{m-1} \tag{8.22}$$

Only  $m$  different coefficients  $C_i$  have to be estimated to determine the whole set of  $n = p^{m-1}$  coefficients  $c_j$ .

Table 8.3 presents such a component scheme for a time series in  $\vartheta$ -time with a basic grid of 30 min,  $p = 4$ , and 7 components ( $m = 7$ ). An interval of 30-min in  $\vartheta$ -time means only some 7 min during the daily volatility peaks around 14:00 GMT, some 80 min during the Far Eastern lunch break, and even more during weekends and holidays with their very low volatility. Table 8.3 shows the minimum relevant time intervals of a component rather than the total size of the volatility memory. In fact, the memory of the volatility can greatly exceed the indicated interval. The medium-term traders of component 5, for example, are not interested in the volatility of hourly returns, are most interested in volatilities observed over 1 to  $3\frac{1}{2}$  days, and are also interested in volatilities observed over longer intervals.

The choice of the number of components,  $m = 7$ , and the factor between the typical time resolutions of the components,  $p = 4$ , is reasonable but somewhat arbitrary. The essential results of this chapter do not strongly depend on this choice and can be found also with other  $m$  and  $p$  values. The model should cover the variety of relevant time resolutions of the market.<sup>9</sup> A too small choice of  $m$  misses the chance of revealing the component structure; a too large  $m$  (with a small  $p$ ) leads to too many parameters to be estimated and an unrealistically narrow definition of market components.

A quantity more suitable for the intuitive understanding than a coefficient  $C_i$  is the impact  $I_i$  of a market component. The expected variance formula, Equation 8.18, strongly suggests a definition of the impact of the  $j^{\text{th}}$  coefficient as  $jc_j$ . The impact  $I_i$  of the  $i^{\text{th}}$  component is defined as the sum of the impacts of all its coefficients  $c_j$ . By inserting the coefficient definitions of Equation 8.22, we obtain

$$I_1 = c_1 = C_1$$

$$I_i = \sum_{j=p^{i-2}+1}^{p^{i-1}} j c_j = (p-1) p^{i-2} \frac{p^{i-1} + p^{i-2} + 1}{2} C_i \quad \text{for } i > 1 \quad (8.23)$$

The impact of the long-term components may be considerable even when the coefficients appear to be small. The impact of the fifth component, for example, is  $I_5 = 30816C_5$ , where  $p$  is assumed to be 4 as in Table 8.3.

The stationarity condition (Equation 8.19) can now be formulated in terms of the impacts. Their sum is smaller than 1:

$$\sum_{i=1}^m I_i < 1 \quad (8.24)$$

<sup>9</sup> The choice of Table 8.3 is  $m = 7$ . An even higher value,  $m = 8$ , has also been tested, leading to a low, insignificant impact of the eighth component and a rejection in a likelihood ratio test. We conclude that the seventh component is the last relevant one on the long-term side.

**TABLE 8.3** Definition of the HARCH components.

A HARCH model with seven market components, each with a range of indices  $j$ . All coefficients  $c_j$  of a component are identical and only seven parameters need to be estimated. The time intervals are the relevant intervals for the volatility perception of the time components (not the total duration of their memory). These basic intervals are given in  $\vartheta$ -time and also in physical time. Two columns show the minimum and maximum size of the interval that can occur for a time component, depending also on the daytime. These time component descriptions may contribute to a better understanding of the model.

$i$	Range of $j$	Approximate range of time intervals in $\vartheta$ -time	Shortest interval (at daily or weekly volatility peaks) in physical time	Longest interval (but avoiding weekends and holidays) in physical time	Description of the time component
1	1	30 min	7 min	80 min	Short-term, intraday dealers (arbitrage opportunities), market makers
2	2 — 4	1 — 2 hr	16 min	3½ hr	Intraday dealers with few transactions per day
3	5 — 16	2½ — 8 hr	50 min	9 hr	Dealers with overnight positions and occasional intraday transactions
4	17 — 64	8½ — 32 hr	4 hr	1 day	Few traders concerned (time intervals often beyond local business hours but less than a business day)
5	65 — 256	1½ — 5 days	1 day	3½ days	Medium-term traders, no intraday trading
6	257 — 1024	5½ — 21 days	3½ days	21 days (weekends always contained)	Long-term traders
7	1025 — 4096	3 — 12 weeks	3 weeks	12 weeks (weekends always contained)	Long-term investors, central banks

TABLE 8.3 Definition of the HARCH components.

### 8.3.3 Generalization of the Process Equation

In Equation 8.13, all the returns considered by the variance equation are observed over “recent” intervals ending at time  $t - 1$ . This strong limitation will be justified by its empirical success, but we can also formulate a more general process equation with observation intervals ending in the past, before  $t - 1$ ,

$$\begin{aligned} r_t &= \sigma_t \varepsilon_t \\ \sigma_t^2 &= c_0 + \sum_{j=1}^n \sum_{k=1}^j c_{jk} \left( \sum_{i=k}^j r_{t-i} \right)^2 + \sum_{i=1}^q b_i \sigma_{t-i}^2 \end{aligned} \quad (8.25)$$

where

$$\begin{aligned} c_0 > 0, \quad c_{jk} \geq 0 & \text{ for } j = 1 \dots n, \quad k = 1 \dots j; \\ b_i \geq 0 & \text{ for } i = 1 \dots q \end{aligned} \quad (8.26)$$

This generalized process equation considers all returns between any pair of two time points in the period between  $t - n$  and  $t - 1$ . It covers the case of HARCH (all  $c_{jk} = 0$  except some  $c_{j1}$ ) as well as that of ARCH and GARCH (all  $c_{jk} = 0$  except some  $c_{jj}$ ). The last term of the variance equation is a “GARCH term,” which may contribute to a parsimonious model formulation. Such a GARCH term may partially model the fading volatility memory of several market components together, but therefore miss the chance of clearly indicating the actual component structure. The main idea of HARCH, taking intervals of different sizes, may also be combined with other ideas from the recent literature about GARCH variations. HARCH can also be regarded as a special case of the Quadratic ARCH model suggested by Sentana (1991). The results obtained for QARCH also apply to HARCH. However, QARCH has been developed in a very different context. Sentana (1991) gives neither a concept of volatilities observed over long intervals nor the stationarity and moment conditions as in Section 8.3.1.

For HARCH, the simple form of Equation 8.13 is preferred. This HARCH is successful in empirical studies, but its computation and estimation is tedious because of the large number of coefficients  $c_j$ . This can be strongly improved by introducing the EMA-HARCH process with its partial volatility concept in the next section.

### 8.3.4 EMA-HARCH Model

In HARCH, the coefficients  $c_1 \dots c_n$  are not regarded as free parameters of the model. The heterogeneous market approach leads to a low number of free model parameters, which determine a much higher number  $n$  of dependent coefficients modeling the long memory of volatility.

The approach is to keep in the equation only a handful of representative interval sizes instead of keeping all of them, and replace the influence of the neighboring interval sizes by an exponential moving average (EMA) of the returns measured on each interval. This also has the advantage of including a memory of the past

intervals. Let us now introduce the concept of *partial* volatility  $\sigma_j^2$ , which can be regarded as the contribution of the  $j^{\text{th}}$  component to the total market volatility  $\sigma^2$ . Here the volatility  $\sigma_j^2$  is defined as the volatility observed over an interval of size  $k_j$ . We can reformulate the HARCH equation in terms of  $\sigma_j$  as follows:

$$r_t = \sigma_t \varepsilon_t$$

$$\sigma_t^2 = c_0 + \sum_{j=1}^n C_j \sigma_{j,t}^2 \quad (8.27)$$

where  $n$  is now the number of time components in the model. The notation is slightly changed to  $C_j$  instead of  $c_j$  used in the old formulation to reflect the different meaning of the coefficients. Unlike the standard HARCH, the partial volatility  $\sigma_j^2$  has a memory of the volatility of *past* intervals of size  $k_j$ . The formal definition of  $\sigma_j^2$  is

$$\sigma_{j,t}^2 = \mu_j \sigma_{j,t-1}^2 + (1 - \mu_j) \left( \sum_{i=1}^{k_j} r_{t-i} \right)^2 \quad (8.28)$$

where  $k_j$  is the aggregation factor of the returns and takes  $n$  possible values, following the relation

$$k_j = p^{j-2} + 1 \quad \text{for } j > 1 \quad \text{with } k_1 \equiv 1 \quad (8.29)$$

When  $p = 4$ ,  $k_j$  can only take the values 1, 2, 5, 17, 65, 257, 1025,  $\dots$ ,  $4^{n-2} + 1$ . For a 5-min data series, this would mean that the horizons would correspond to 5 min, 10 min, 25 min and so on. The construction of Equation 8.29 ensures that the time components ( $k_j$ 's) are economically meaningful. Equation 8.28 is the iterative formula for an exponentially weighted moving average, a special application of Equation 3.51. The volatility memory is defined as a moving average of recent volatility. The depth of the volatility memory is determined by the constant  $\mu_j$ :

$$\mu_j = e^{-\frac{1}{M(k_j)}} \quad (8.30)$$

where the memory decay time constant of the component is given as the function  $M$  of the component's volatility interval  $k_j$ . Without introducing new parameters,  $M(k_j)$  can be defined as

$$M(k_j) = \frac{(k_{j+1} - k_j)}{2} \quad (8.31)$$

This definition is based on the start and the end point of the component interval  $k_j$  and makes sure that the EMA kernel is centered at the characteristic time horizon of the component.



It is easy to prove that a necessary stationarity condition for the new formulation is

$$\sum_{j=1}^n k_j C_j < 1 \quad (8.32)$$

The proof relies on the fact that the expectation of the exponential moving average is the same as the expectation of the underlying time series and that the expectation of cross terms is zero. A similar proof as in Dacorogna *et al.* (1998a) can be given for the sufficiency of this condition.

We can now define the impact  $I_j$  of each component,

$$I_j = k_j C_j \quad (8.33)$$

Every component with a coefficient  $C_j$  has an impact  $I_j$  on the conditional volatility process. The stationarity condition, Equation 8.32, can be re-formulated using the sum of impacts:

$$\sum_{j=1}^n I_j < 1 \quad (8.34)$$

An iterative formula needs an initial value for  $\sigma_j^2$  at the beginning of the time series. A reasonable assumption of that initial value is the unconditional expectation of  $\sigma_{j,t}^2$ . Here the first value is computed from a data sample prior to the model estimation sample. We term this sample the “build-up” sample.

### 8.3.5 Estimating HARCH and EMA-HARCH Models

HARCH and EMA-HARCH models are applied to and estimated for empirical FX data here. The time series are homogeneous in  $\vartheta$ -time, which removes the seasonal pattern of intraday volatility. The time interval is 30 min and the sample includes 10 years of data from January 1, 1987, to December 31, 1996.

For the computation and estimation of both HARCH and EMA-HARCH, we use seven components. For HARCH, the components of Table 8.3 are used. For EMA-HARCH, one component is built from only one time interval but includes, according to Equation 8.28, a moving average that extends over a certain range, which should account for the neighboring time intervals. In fact, there are now two parameters controlling the component definition. The time interval size over which returns are computed,  $k_j$ , and the range of the moving average,  $M(k_j)$ . Both of them are fixed and the optimization is carried out to solve the  $C_j$  parameters. The optimization is implemented by searching for the maximum of the log-likelihood function. The procedure we follow to find the maximum is described at the end of Section 8.2.1. It combines two methods: a genetic algorithm (GA) search (Pictet *et al.*, 1995) and the Berndt, Hall, Hall, and Hausman (BHHH) algorithm (Berndt *et al.*, 1974).

TABLE 8.3: HARCH and EMA-HARCH Models

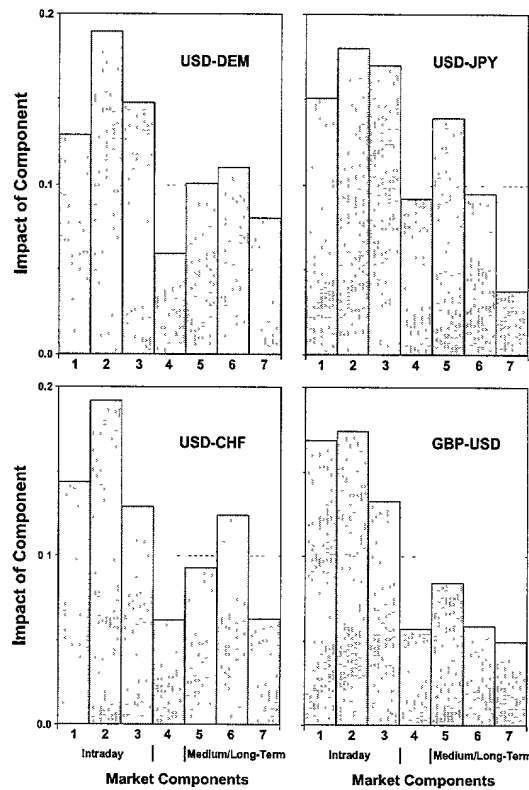
**TABLE 8.4** HARCH coefficients for USD-DEM.

Comparison between the coefficients and impacts of the two HARCH processes from a half-hourly USD-DEM series, which is equally spaced in  $\vartheta$ -time over 10 years. Instead of the coefficients  $C_i$ , the impacts  $I_i$  are given. These provide a direct measure of the impacts of the market components on the HARCH variance. The market components are those of Table 8.3 for HARCH and as in Equations 8.28 and 8.30 for EMA-HARCH. The distribution of the random variable  $\varepsilon(t)$  is normal with zero mean and a unit variance.

USD-DEM Coefficient	HARCH			EMA-HARCH		
	Estimate	Standard error	<i>t</i> -statistics	Estimate	Standard error	<i>t</i> -statistics
$c_0$	$1.276 \times 10^{-7}$	$0.03994 \times 10^{-7}$	31.94	$0.529 \times 10^{-7}$	$0.04399 \times 10^{-7}$	21.01
$I_1$	0.1309	0.007151	18.30	0.1476	0.008295	17.80
$I_2$	0.1930	0.010010	19.28	0.1875	0.012297	15.25
$I_3$	0.1618	0.009179	17.62	0.1829	0.012545	14.58
$I_4$	0.0703	0.007363	9.55	0.0507	0.010324	4.91
$I_5$	0.1003	0.006774	14.81	0.1434	0.010952	13.10
$I_6$	0.1014	0.006892	14.71	0.1120	0.011835	9.47
$I_7$	0.0990	0.006118	16.18	0.1145	0.010540	10.86
Log-likelihood		5.7947			5.8014	

The result of the optimization procedure is a set of  $C_j$  coefficients from which the component impacts are calculated using Equation 8.23 (for HARCH) or Equation 8.33 (for EMA-HARCH). The sum of impacts  $I_j$  must be below one for stationarity of the process (Equations 8.24 and 8.34). In Table 8.4, the coefficients for both the HARCH and EMA-HARCH are shown with their *t*-statistics for USD-DEM. They are obtained on exactly the same data set. The log-likelihoods can be compared because both models have the same number of independent coefficients. Clearly, the log-likelihood is improved by going to EMA-HARCH. In both cases, all coefficients are highly significant according to the *t*-statistics and contribute to the variance equation. The stationarity property is fulfilled in both cases. The HARCH and EMA-HARCH have total impacts of 0.8567 and 0.9386, respectively. The impacts of the different components are remarkably similar. Two small differences are worth noticing. The relative importance of the long-term components is slightly higher for EMA-HARCH (37% instead of 35%) and the minimum for the fourth component is more pronounced in EMA-HARCH. The *t*-statistics are also consistently smaller for EMA-HARCH than for HARCH but still highly significant in all cases. The residuals in both formulations still present an excess kurtosis, as was noticed in Müller *et al.* (1997a) for HARCH.

These results show that we have achieved the goal of redesigning the HARCH process in terms of moving averages. We are able to keep and even improve on the properties of the original HARCH and to considerably reduce the computational



**FIGURE 8.5** Impacts of market components of HARCH processes with components as defined in Table 8.3. Each HARCH model has been made for a particular FX rate by fitting a half-hourly time series equally spaced in  $\vartheta$ -time over 7 years. The differences between the impacts, in particular the low values of the fourth component, are highly significant (see the error values of Table 8.4). The values for USD-DEM are those presented in Table 8.4 and they are not fundamentally different from those of other FX rates.

time to optimize the model. The EMA formulation of the process equation reduces this time by a factor of 1000, making the computation of HARCH volatility much more tractable even with limited computational power. In the next section, we will explore the forecasting ability of these models and compare it to a more traditional approach to volatility.

The impacts  $I_i$  are also plotted in a histogram (Figure 8.5) where it is possible to compare the results for different FX rates. The impact of the fourth component is the weakest among all impacts. This is not only for USD-DEM but also many other rates and also for other sampling periods.<sup>10</sup> The fourth component has a

<sup>10</sup> When a 7-year sample is split into two parts of 3 1/2 years, the estimated coefficients on both of these subsamples are quite stable.

**TABLE 8.5** Results of the EMA-HARCH for the LIFFE Three-Month Euromark.

Results of the EMA-HARCH process estimate for 3-hr  $\vartheta$ -time intervals for the different forward rates for the LIFFE Three-Month Euromark. The underlying data are from the LIFFE Three-Month Euromark. Standard errors are given. Instead of the coefficients  $C_i$  (for  $i > 0$ ), the corresponding impacts  $I_i$  are given. Data sample: from April 6, 1992, to December 30, 1997, representing 16,774 observations. The forward rates are labeled according to the market conventions for forward rate agreements, as explained in the text. The 3x6 forward interest rate, for example, applies to the interval starting in 3 months and ending in 6 months.

	3x6	6x9	9x12	12x15	15x18
$K_0$	$2.90 \pm 0.15$	$4.54 \pm 0.24$	$4.29 \pm 0.24$	$6.45 \pm 0.39$	$4.22 \pm 0.37$
$I_1$	$0.20 \pm 0.01$	$0.17 \pm 0.01$	$0.19 \pm 0.01$	$0.18 \pm 0.01$	$0.12 \pm 0.01$
$I_2$	$0.01 \pm 0.01$	$0.00 \pm 0.02$	$0.00 \pm 0.02$	$0.01 \pm 0.01$	$0.00 \pm 0.02$
$I_3$	$0.17 \pm 0.02$	$0.16 \pm 0.02$	$0.16 \pm 0.02$	$0.15 \pm 0.02$	$0.15 \pm 0.02$
$I_4$	$0.08 \pm 0.02$	$0.11 \pm 0.02$	$0.13 \pm 0.02$	$0.14 \pm 0.02$	$0.12 \pm 0.02$
$I_5$	$0.11 \pm 0.02$	$0.15 \pm 0.02$	$0.16 \pm 0.02$	$0.15 \pm 0.02$	$0.19 \pm 0.02$
$I_6$	$0.23 \pm 0.02$	$0.22 \pm 0.02$	$0.19 \pm 0.02$	$0.08 \pm 0.02$	$0.11 \pm 0.02$
$I_7$	$0.00 \pm 0.01$	$0.03 \pm 0.01$	$0.04 \pm 0.01$	$0.02 \pm 0.01$	$0.03 \pm 0.01$
$L$	7.753	7.478	7.345	7.320	7.307

typical time horizon of around 12 hr—too long for intraday dealers and too short for other traders. This naturally explains the weakness of that component.

When comparing the impacts of Figure 8.5 with the component definition of Table 8.3, we see further interesting features captured by HARCH models. First, the short-term components have, in all cases, the largest impacts. These short-term components model essentially the intraday dealers and the market makers who are known to dominate the market. Second, the similarity in the impacts of the USD-CHF and USD-DEM are plausible as it is well known that the Swiss National Bank policy was tightly tied to the USD-CHF to the USD-DEM rates. The relative weakness of the longer-term components for the GBP-USD is another relevant piece of information that can be gathered from this parametrization and has been confirmed to us by market participants. Since the 1992 crisis, the long-term investors have been reluctant to invest in this market and have been more concentrated on the cross rate GBP-DEM. The relative impact of the fifth and the sixth components are in the same order for USD-CHF and USD-DEM but inverted in the case of both USD-JPY and GBP-USD.

### 8.3.6 HARCH in Interest Rate Modeling

As described in Chapter 7, we have performed a lead-lag correlation analysis and established the HARCH effect for forward interest rates implied by interest rate futures, constructed according to Section 2.4.2 (see Figure 7.9). In this section, we use the HARCH parametrization in terms of market components to investigate

TABLE 8.5 "continued"

whether we obtain similar features as in the case of foreign exchange rates. To avoid a systematic, deterministic decrease of volatility as explained by Section 5.6.4, we use forward rates for fixed time intervals. The forward rates are labeled according to the market conventions for forward rate agreements. The  $I \times J$  forward rate (e.g., the 3x6 forward rate) is the forward rate quoted at time  $t$  and applicable for the interval starting at time  $(t + I)$  and ending at time  $(t + J)$  ( $I$  and  $J$  are expressed in months). The corresponding time-to-start is  $I$  months and the maturity is  $(J - I)$  months.

The results of the EMA-HARCH process estimation for 3-hr  $\vartheta$ -time intervals for the different forward rates for the LIFFE Three-Month Euromark are given in Table 8.5. The EMA-HARCH process is estimated for each forward rate by maximizing the log-likelihood and use data from April 6, 1992, to December 30, 1997, with about 16,800 observations. The distribution of the random variable  $\varepsilon_t$  is normal with zero mean and a unit variance. The market components (with  $p = 4$ ) are similar to the ones described in the previous section. Like in the case of FX rates, the impact coefficient for the interval range from 6 hr to 1 days (second component) is very small. These results also indicate a decreasing impact of the longer-term components (corresponding to the market actors with the longest time-horizon) going from the first forward rate (i.e., whose time-to-start is closest in the future) to the last one (i.e., whose time-to-start is farthest in the future), reflecting the decrease in the volatility autocorrelation. The sum of the impacts is smaller than one in all cases, meaning that the estimated processes are stationary according to Equation 8.34.

## 8.4 FORECASTING SHORT-TERM VOLATILITY

The true test of the veracity of a volatility model is its ability to forecast the future behavior of volatility. This means that the data used to test the model are distinct from the data used to find the model parameters. All of the analyses described in this section are performed in an out-of-sample setting.

There is some added complexity in the case of volatility models where there is no unique definition of volatility. Andersen and Bollerslev (1998a) showed that, if the wrong estimators of volatility are taken, it is not possible to really test the forecasting quality of a model. That is why it is important to set a framework in which a forecasting performance analysis can be performed.

### 8.4.1 A Framework to Measure the Forecasting Performance

We choose here a path similar to that proposed in Taylor and Xu (1997). We construct a time series of realized hourly volatility,  $v_h(t)$ , from our time series of returns as follows:

$$v_{h,t} = \sum_{i=1}^{a_h} r_{t-i}^2 \quad (8.35)$$

where  $a_h$  is the aggregation factor. In this case, we use data points every 10 min in  $\vartheta$ -time, so the aggregation factor is  $a_h = 6$ .

Forecasts of different models are compared to the realized volatility of Equation 8.35. The one-step ahead forecasts are based on hourly returns in  $\vartheta$ -time. The advantage of using hourly returns instead of 30-min returns as in the previous section is that hourly forecasts are compatible with the historical hourly volatility defined in Equation 8.35. Four models are studied here.

- The first model is used as a benchmark and is a naive historical model inspired by the effect described in Section 7.4.2 and Müller *et al.* (1997a) where low-frequency volatility predicts high-frequency volatility. The historical volatility is computed over a given day measured from the hourly returns. This quantity, properly normalized, is used as a predictor for the next hour volatility,  $v(t + 1)$ , as defined in Equation 8.35. Formally the forecasting model  $v_b$  is

$$v_{b,t} = \frac{1}{24} \sum_{j=1}^{24} \left( \sum_{i=6(j-1)+1}^{6j} r_{t-i} \right)^2 \quad (8.36)$$

where the factor in front of the summation is here to normalize  $v_b$  to hourly volatility.

- GARCH(1,1) is

$$v_{garch,t} = h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1} \quad (8.37)$$

where  $\varepsilon_t$  is i.i.d. and follows a normal distribution function with zero mean and unit variance.

- The HARCH model in Equation 8.13 and the seven components of Table 8.3, introduced by Müller *et al.* (1997a).
- The EMA-HARCH model in Equation 8.27 and 8.28 with seven components.

The three parameter models are optimized over a sample of 5 years of hourly data using the estimation procedure described in Section 9.3.5. The forecasts are then analyzed over the 5 remaining years. We term this procedure the static optimization. To account for possible changes in the model parameters, we also recompute them every year using a moving sample of 5 years. We term this procedure dynamic optimization. In this case, the performance is always tested outside of the gliding sample to ensure that the test is fully out-of-sample. In both cases, we use an out-of-sample period of 5 years of hourly data, which represents more than 43,000 observations.

We compare the accuracy of the four forecasting models to the realized hourly volatility of Equation 8.35. The quantities of interest are the forecasting signal,

$$s_f = \tilde{v}_{f,t} - v_{h,t} \quad (8.38)$$

where  $\tilde{v}_f$  is any of the forecasting models, and the realized signal,

$$s_r = v_{h,t+1} - v_{h,t} \quad (8.39)$$

The quantity  $\tilde{v}_{f,t}$  is taken in the estimation sample either directly or rescaled by the ratio of the averages  $\bar{v}_h$  and  $\bar{\tilde{v}}_f$ . This makes the forecast values on average closer to the historical volatility. In the rest of this chapter, we call the quantity,  $(\bar{v}_h \cdot \tilde{v}_{f,t})/\bar{\tilde{v}}_f$ , the rescaled forecast.

In this formulation, performance measures proposed in Dacorogna *et al.* (1996) can be applied because the quantities defined in Equations 8.38 and 8.39 can take positive and negative values contrary to the volatilities, which are positive definite quantities. One of these measures is the direction quality,

$$Q_d = \frac{\mathcal{N}(\{\tilde{v}_f | s_f \cdot s_r > 0\})}{\mathcal{N}(\{\tilde{v}_f | s_f \cdot s_r \neq 0\})} \quad (8.40)$$

where  $\mathcal{N}$  is a function that gives the number of elements of a particular set of variables. It should be noted that this definition does not test the cases where either the forecast is the same as the current volatility or the volatility at time  $t + 1$  is the same as the current one. This is, of course, unlikely to occur in our particular case. A detailed statistical discussion of this measure can be found in Pesaran and Timmerman (1992).

In addition to this measure, we use a measure that combines the size of the movements and the direction quality. It is often called the *realized potential*,

$$Q_r = \frac{\sum \text{sign}(s_f \cdot s_r) |s_r|}{\sum |s_r|} \quad (8.41)$$

In fact, the measures  $Q_r$  and  $Q_d$  are not independent and  $Q_r$  is a weighted average of  $\text{sign}(s_f \cdot s_r)$  whereas  $2Q_d - 1$  is the corresponding unweighted average. It is easy to show that if

$$Q_r > 2Q_d - 1 \quad (8.42)$$

the forecast of the sign of  $s_r$  for large  $|s_r|$  values is better than average.

A more traditional measure such as the comparison of the absolute error of a model to a benchmark model can also be used. This benchmark model is chosen to be the historical volatility as defined in Equation 8.36,  $v_b$ . The quantity

$$Q_f = 1 - \frac{\sum |s_r - s_f^{model}|}{\sum |s_r - s_f^{benchmark}|} \quad (8.43)$$

is a quality measure, which increases with an increasing performance of the model. If  $Q_f > 0$ , the model outperforms the benchmark. If  $Q_f < 0$ , the benchmark outperforms the model. The second part of Equation 8.43 is similar to the known Theil's  $U$ -statistic, Makridakis *et al.* (1983), except that we use the absolute value instead of the squared errors.

**TABLE 8.6** Forecasting performance for USD-DEM.

Forecasting accuracy of various models in predicting short-term market volatility. The performance is measured every hour over 5 years, from January 1, 1992, to December 31, 1996, with 43,230 observations. In parentheses, the accuracy of rescaled forecasts is shown.

USD-DEM	$Q_d$	$Q_r$	$Q_f$
Static Optimization			
Benchmark	67.7% (67.6%)	54.2% (54.3%)	0.000
GARCH(1,1)	67.8% (67.3%)	58.5% (59.7%)	0.085 (0.072)
HARCH(7c)	69.2% (68.7%)	58.3% (59.2%)	0.134 (0.129)
EMA-HARCH(7)	69.4% (68.8%)	60.7% (62.5%)	0.140 (0.128)
Dynamic Optimization			
Benchmark	67.7% (67.4%)	54.2% (54.6%)	0.000
GARCH(1,1)	67.0% (66.0%)	59.5% (59.8%)	0.074 (0.057)
HARCH(7c)	67.7% (66.8%)	60.1% (60.8%)	0.113 (0.102)
EMA-HARCH(7)	68.8% (67.7%)	62.4% (62.9%)	0.133 (0.117)

The summations (including  $N$ ) in Equations 8.40, 8.41, and 8.43 are over all hours in the out-of-sample period. The number of independent observations is large so that the degrees of freedom of the calculated tests are sufficiently large. Performance measures based on squares such as the signal correlation or squared errors are not used because our interest is in squared returns and the fourth moment of the distribution of returns may not be finite, as discussed in Section 5.4.2.

#### 8.4.2 Performance of ARCH-Type Models

In Table 8.6, the results for the different performance measures are presented for the most traded FX rate, USD-DEM, for the static and dynamic optimizations. In parentheses, the results for the scaled forecasts are presented. For all measures, three parameter models perform better than the benchmark and the EMA-HARCH performs the best. The forecast accuracy is remarkable for all ARCH-type models. In more than two-thirds of the cases, the forecast direction is correctly predicted and the mean absolute errors are smaller than the benchmark errors for all models. The realized potential measure shows that the forecast of volatility change is accurate not only for small  $|s_r|$  but also for large ones. The condition expressed in Equation 8.42 is always satisfied for all models. Neither the scaled forecast nor the dynamic optimization seems to significantly improve the forecasting accuracy. The realized potential  $Q_r$  is the only measure that consistently improves with



dynamic optimization. Examining the model coefficients computed in moving samples shows that they oscillate around mean values. No structural changes in the coefficients were detected. The accuracy improvement in  $Q_r$  together with the loss in  $Q_f$  in the case of dynamic optimization indicates that the prediction of large movements is improved at the cost of the prediction of direction of small real movements. From the point of view of forecasting short-term volatility, the EMA-HARCH is the best of the models considered here and compares favorably to HARCH. Similar conclusions can be drawn from the results for four other FX rates.<sup>11</sup> The cross rate JPY-DEM presents results slightly less accurate than the other currencies, but it should be noted that the early half of the sample has been synthetically computed from USD-DEM and USD-JPY. This may lead to noise in the computation of hourly volatility and affect the forecast quality.

<sup>11</sup> The interested reader will find them in Dacorogna *et al.* (1998b), where similar tables are listed for USD-JPY, GBP-USD, USD-CHF, and DEM-JPY.

# 9

---

## FORECASTING RISK AND RETURN

### 9.1 INTRODUCTION TO FORECASTING

This section examines forecasting models for different variables. The predicted variable should be observable, so the forecasts and the true variable values can be compared in the future to allow for statistical forecast quality tests. The following variables can be predicted:

- The absolute size of future returns. This can be done in different mathematical forms, one of them using *realized volatility* as defined in Section 3.2.4. We ignore the direction of the future price returns here and assume their probability distribution function to be symmetric by default. Under this assumption, the chosen forecast variable also measures the *risk* of holding a position in the financial asset.
- The future return over the forecast period, including its sign. This also implies a price forecast because the future price is the price now plus the future return.
- A full probability distribution function of future returns. This is the most comprehensive goal. Given the natural uncertainty of forecasting, we are rarely able or willing to forecast details of the distribution function, and we are more than happy to have good forecasts of its center and its width.

There are basically two approaches for constructing forecasting models. The foreign exchange market again serves as our main example. The first approach builds upon structural economic models testing various forms of market efficiency<sup>1</sup> or the study of the issues such as the purchasing power parity model and the modeling of risk premia (see Baillie and McMahon, 1989; MacDonald and Taylor, 1992). Meese and Rogoff (1983) carried out the first comprehensive out-of-sample tests of these models, which they call structural models.

Models following the second approach are often called time series models and are based on information extracted from the past of the time series through various forms of linear and nonlinear statistical operators and prefiltering techniques. These types of models can be univariate or multivariate.<sup>2</sup> In this chapter, we adopt the second approach and study univariate time series models by utilizing only past prices to forecast a given series. There are two main motivations for this approach. First, the absence of any theory for the short-term movements of the foreign exchange (FX) rates makes the structural models irrelevant for these horizons. Second, the availability of high-frequency data can capture many of the market effects that are relevant to the short-term movements, (e.g., the behavior of different market participants).

The forecasting models presented in this chapter are *univariate* where only one time series is predicted. They are univariate not only in the predicted target variable but also in the information set used. Multivariate forecasting as an important but complex subject is not discussed here, but Chapter 10 has some relevant discussions.

The models work with high-frequency data as described by Chapter 2 and take into account every tick in the market. The predicted quantity (e.g., the price or future realized volatility) is related to a *time horizon*, (e.g., the return of the next hour or the volatility of one full business day from now.) The use of high-frequency data allows us to make short-term forecasts for time intervals less than a day. This leads to a large number of observed forecast intervals and thus high statistical significance.

In principle, the knowledge of the “true” data-generating process in the sense of Chapter 8 should also lead to the “true” forecasting model. We have indeed used statistical processes to generate forecasts and measured the success of these statistical processes in terms of their forecasting quality in Chapter 8. In practice, the way from a statistical process to a good forecasting model is not as straightforward. Many otherwise popular statistical processes have serious shortcomings when looking at the intradaily and temporally aggregated behavior, as shown in Section 8.2. Moreover, the statistical processes of that section are volatility models. The price aspect of these models is trivial by having the current price as expectation value for future prices. When moving to forecasting models, we can be more ambitious by also constructing nontrivial price forecasts. We also introduce new

<sup>1</sup> The reader may refer to Fama (1970, 1991).

<sup>2</sup> Granger and Newbold (1977) and Priestley (1989) are introductions to these types of models.

testing methods for forecasts. Thus the two topics, data-generating processes and forecasting, are only loosely related.

Forecasting models can be tested by comparing the forecasts to the actual values of the predicted variable. A possible test criterion is the standard deviation of the forecasts from the actual values. Different test criteria can be computed by statistical means, using a test data sample as discussed in Section 9.4. The test result consists of not only a quantitative quality measure but also a statistical significance measure of this quality. The test sample can also be used to optimize the forecasting model and its parameters. In that case, the resulting optimized model should be tested in *another* sample (i.e., out-of-sample). The test results of the original sample (in-sample) cannot be used as an unbiased measure; they only give an upper limit of the otherwise unknown forecasting model quality.

Two examples of univariate time series models are given: volatility forecasting models used for risk assessment in Section 9.2 and a large real-time price forecasting system with live data feeds in Sections 9.3 and 9.4.3.<sup>3</sup>

## 9.2 FORECASTING VOLATILITY FOR VALUE-AT-RISK

Risk can be measured by different means, for example, through an extreme value analysis as in Sections 5.4.2 and 5.4.3. Here we follow a simpler approach by regarding volatility as the variable that determines the risk. This is also the view of popular risk assessment methods. In these methods, the volatility value is inserted in a standard model to compute the Value-at-Risk (VaR): the expected loss of a portfolio after one business day corresponding to the 1% quantile,<sup>4</sup> (i.e., in a scenario that is worse than 99% of the expected cases and better than the remaining 1%). Inserting a volatility figure (computed from variances and covariances of returns of the portfolio assets) may not be enough to compute a reliable VaR. This is discussed by Davé and Stahl (1998), but is not the focus of interest here.

The required volatility value is in fact a volatility *forecast* for the period from “now” to “now plus one business day.” In this section, we discuss univariate volatility forecasting models. Multivariate volatility models need separate treatment because they depend on the intradaily covariance or correlation between assets. This poses some problems as discussed in Chapter 10.

### 9.2.1 Three Simple Volatility Forecasting Models

Müller (2000) has a discussion of volatility forecasts based on time series operators as presented in Section 3.3. Following that paper, we consider three operator-based volatility forecasting methods of increasing sophistication and quality: (1) the volatility forecasts of RiskMetrics™ developed by J. P. Morgan (1996) as a

<sup>3</sup> This forecasting model is running in real time as a part of the Olsen & Associates Information System (OIS).

<sup>4</sup> Our scientific interest also extends to forecast intervals other than one business day and quantiles other than 1%, of course.

well-known example, (2) an improved version based on tick-by-tick data, and (3) a further improved multi-horizon version.

All of these volatility models can be seen as observations of volatility in the past (i.e., realized volatility measurements as Equations 3.8 or 3.68, for example). However, the models are intended to be applied to the future. The computed volatility values, although measured in the past, are estimates of the *future* volatility and thus measures of risk. Autoregressive heteroskedasticity as discussed in Section 5.6.1 is the stylized fact that justifies using a certain past volatility as an estimate of future volatility. Section 9.3 has another approach where the volatility forecast is no longer a realized volatility of the past, and Section 8.4.2 considers volatility forecasts directly derived from statistical processes.

The RiskMetrics methodology uses a well-known example of a simple volatility forecast based on an IGARCH process with the following conditional expectation of the squared return:

$$\sigma^2(t) = \mu \sigma^2(t - \Delta t) + (1 - \mu) [x(t) - x(t - \Delta t)]^2 \quad (9.1)$$

with  $\mu = 0.94$ . This formula is evaluated only once per business day, at a given daytime; the resulting volatility value is valid until it is replaced by a new one, one business day later. The time scale  $t$  is thus a business time scale omitting weekends, with  $\Delta t = 1$  business day. Equation 9.1 is an exponential moving average (EMA) iteration as explained in Section 3.3.5 and can be written as such, using the notation of Equation 3.51,

$$\sigma^2(t) = \text{EMA} \left[ \tau; [x(t) - x(t - \Delta t)]^2 \right] \quad (9.2)$$

evaluated at discrete time points separated by  $\Delta t = 1$  business day, with an EMA range of  $\tau = \mu/(1 - \mu) = 15.67$  business days. The EMA operator is explained in Section 3.3.5, but Equation 3.52 has to be replaced here by a version for discrete, homogeneous time series,


$$\mu = \nu = \frac{1}{1 + \alpha} = \frac{\tau}{\tau + 1} \quad (9.3)$$

as explained by Müller (1991). The only parameter,  $\mu = 0.94$ , has been chosen to optimize the volatility forecasting quality of Equation 9.1 over a wide range of financial assets and test periods according to J. P. Morgan (1996).

In Figure 9.1, two volatilities are presented. The difference between the two curves solely originates from the choice of daytime when the price  $x$  is sampled and the volatility is computed by Equation 9.1 or 9.2. One curve is sampled at 7 a.m. (Greenwich Mean Time) GMT which is a time in the late afternoon of East Asian time zones or a suitable daytime for the daily risk calculations of an East Asian risk manager. The other curve is sampled at 5 p.m. GMT, a suitable daytime for a risk manager in London.

The differences between the two curves are surprisingly large: up to 25%, an alarming uncertainty for risk managers. In our case, two risk managers measure

FOOTNOTES



**FIGURE 9.1** Top panel: Daily standard RiskMetrics USD-JPY volatility for January, 1999 to February 1999. Circles: Data sampled at 7 a.m. GMT. Diamonds: Data sampled at 5 p.m. GMT. Bottom panel: The USD-JPY price plotted against time.

very different volatility and thus risk levels for the *same* financial instrument, just because they live in different time zones. A difference can persist over weeks, as shown in Figure 9.1. This figure is just an example. The same surprisingly strong effect can also be found for other financial instruments, sampling periods, choices of daytime, and process equations.

Both deviating volatility values cannot be right at the same time; there must be an error in these values. This error is of stochastic nature; there is no systematic bias dependent on the daytime. In Figure 9.1, the difference between the two curves is neither always positive nor negative; it changes its sign.

Figure 9.1 demonstrates the large stochastic error of the RiskMetrics method. The large size of this error has two main reasons:

1. The rather small range of the kernel of about 16 business days. The number of independent observations is limited. We cannot essentially change this fact, because the choice of a short range is also motivated by the goal of fast adaptivity to new market events.

2. The results depend on only one observation per day, taken at a certain daytime. All the other information on prices of the day is thrown away. The value at that daytime may have little representation for the full day: it may be located on top of a short-lived local peak of the price curve.

The second investigated volatility forecasting model was introduced by Müller (2000). It follows RiskMetrics as closely as possible. There are only two innovative modifications:

- The squared volatility  $\sigma^2(t)$  is computed at every available tick, not only once per business day.
- Simple returns are replaced by operator-based, smoothed returns.

Nothing is changed otherwise; the sampling range of 15.67 business days and the business-daily nature of (smoothed) returns are preserved. The formula is again written with the help of time series operators:

$$\sigma^2 = c \text{EMA} \left[ \tau; (x - \text{EMA}[\Delta t, 4; x])^2 \right] \quad (9.4)$$

again with  $\Delta t = 1$  business day and  $\tau = 15.67$  business days. Equation 9.4 is iteratively evaluated tick by tick. The iterated operator  $\text{EMA}[\tau, 4; x]$  is defined by Equation 3.53. As the simple EMA operator, it can be efficiently computed by using the iterative Equation 3.51.

The constant  $c$  compensates for the fact that we use smoothed returns,  $x - \text{EMA}[\Delta t, 4; x]$ , instead of the simple returns,  $x(t) - x(t - \Delta t)$ . In the case of  $x$  following a Gaussian random walk, the theoretically correct value is  $c = 128/93$ . Using this factor eliminates a systematic bias of the tick-by-tick volatility as compared to the RiskMetrics volatility.

Equation 9.4 is computed on a special business time scale defined as follows. The 49 weekend hours from Friday 8 p.m. GMT to Sunday 9 p.m. GMT are compressed to the equivalent of only 1 hr outside the weekend. This fully corresponds to the time scale of RiskMetrics, which omits the weekend days. A more sophisticated and appropriate choice of the business time scale would be the  $\vartheta$ -time of Chapter 6, but this is avoided here in order to keep the approach as close to RiskMetrics as possible.

The advantages of the tick-by-tick volatility forecast are demonstrated in Figure 1.3. The volatility as a function of time appears as one continuous, consistent curve. We obtain volatility values at any daytime now, not just once or twice a day. A risk manager in London measures the risk of the instrument on the same basis as a risk manager in East Asia, as should be expected. The variations of the volatility level over time are more moderate in Figure 1.3 than the corresponding variations of the RiskMetrics volatility, although the kernel range of 15.67 business days is the same.

The tick-by-tick volatility forecast is based on (almost) continuously overlapping returns. Overlapping returns lead to reduced stochastic noise of volatility measurements, as shown in Section 3.2.8. In addition to this, the tick-by-tick

volatility is based on smoothed rather than simple returns, which also leads to a reduction of stochastic noise.

The third volatility model is a multiple-horizon version of the second model:

$$\sigma^2 = \frac{\sum_{k=0}^{n-1} f_w^k \sigma_k^2}{\sum_{k=0}^{n-1} f_w^k} \quad (9.5)$$

with

$$\sigma_k^2 = c \text{EMA} \left[ \tau_0 f_\tau^k; \left( x - \text{EMA}[\Delta t_0 f_{\Delta t}^k, 4; x] \right)^2 \right] \quad (9.6)$$

where the partial volatility forecasts of Equation 9.6 are like the volatility forecasts of Equation 9.4. The weights  $f_w^k$  of the partial volatility forecasts, their return intervals  $\Delta t_0 f_{\Delta t}^k$ , and their sampling ranges  $\tau_0 f_\tau^k$  are in geometric sequences and can be flexibly chosen and optimized by setting the parameters  $n$  (the number of partial forecasts),  $f_w$ ,  $\Delta t_0$ ,  $f_{\Delta t}$ ,  $\tau_0$ , and  $f_\tau$ .

The third volatility model (Equation 9.5) shares the advantages of the second one (Equation 9.4) and has the additional multiple-horizon property, which leads to superior volatility forecast quality. This is in analogy to the multiple-horizon EMA-HARCH process shown in Section 8.3.4, which is also superior to single-horizon processes such as GARCH.

### 9.2.2 Choosing the Best Volatility Forecasting Model

The quality of volatility forecasting models has to be measured in statistical tests, comparing the forecasts to the actual values of the target variable, which is a form of realized volatility here.

In out-of-sample tests of the three volatility forecasting models presented in Section 9.2.1, the tick-by-tick model of Equation 9.4 has distinctly better volatility forecasts than the RiskMetrics model of Equation 9.1 or 9.2. Equation 9.5 leads to even better volatility forecasts.

Testing the quality of volatility forecasts implies some technical difficulties. First, there are several quality measures to choose from. This is discussed in Sections 9.4.1 and 8.4.1, where volatility forecasts are derived from process equations and tested by several criteria.

A second difficulty lies in the bias of both realized volatility (the target variable) and volatility forecasts which appears if the return intervals chosen are too small. This bias is discussed in Section 3.2.4 and in Andersen *et al.* (2000). Volatility forecast tests are affected by this bias. A treatment of the bias is almost inevitable when designing volatility forecasting models and tests based on high-frequency returns over intervals of less than an hour. Corsi *et al.* (2001) propose a suitable bias correction method.

Due to these technical difficulties, there is no comprehensive study of high-frequency volatility forecasts and their qualities yet. The final goal is the development of a consistent methodology of risk analysis based on high-frequency data with superior forecasting quality: *real-time risk assessment*.



### 9.3 FORECASTING RETURNS OVER MULTIPLE TIME HORIZONS

This section examines the forecasting model of Dacorogna *et al.* (1996). This model supports several forecast intervals. Hourly returns are predicted as well as daily, weekly, monthly, and quarterly returns. The forecasting model of returns relies on an underlying volatility forecast. Both the volatility and return forecasts use the same methodology. Volatility is treated with the help of an alternative time scale, the *intrinsic time* of the time series.

#### 9.3.1 Intrinsic Time

The foreign exchange returns exhibit conditional heteroskedasticity which can be treated through a change of time scale. This is the second layer of our forecasting model on top of the business time  $\vartheta$ -scale. Some literature followed a similar approach to treat the conditional heteroskedasticity, such as Stock (1988), who uses two types of time deformation, one based on the time series itself and one on business cycle variables.<sup>5</sup> In our approach, we also use the underlying time series to construct a time deformation. It is based on the scaling law defined in Equation 6.2 and on the price volatility:

$$\tau(t_c) \equiv \tau(t_{c-1}) + k \frac{\vartheta(t_c) - \vartheta(t_{c-1})}{\Delta\vartheta} \frac{|\Delta x|^E}{c} \quad (9.7)$$

where  $t_c$  is the current time, the price difference is taken on the same interval as  $\Delta\vartheta$ , and  $E$  and  $c$  are the scaling law inverse exponent and factor, respectively. The constant  $k$  is a calibration factor dependent on the particular time series. Its role is to keep  $\tau$  in line with physical time in the long run. This relationship is in fact the reverse of the scaling law for a particular return taken on a constant  $\vartheta$ -time interval size.

This second new time scale, the  $\tau$ -scale, does not directly use the physical time  $t$ , and does not need to have fundamental information about the behavior of the series. The only information needed to define the scale are the values of the time series themselves. Thus we have chosen to call this time scale *intrinsic time*. The consequence of using such a scale is to expand periods of high volatility and contract those of low volatility, thus better capturing the relative importance of events to the market. Any moving average based on the intrinsic time  $\tau$  dynamically adapts its range to market events. Therefore a forecasting model based on the  $\tau$ -scale has a *dynamic memory* of the price history.

There is, however, a problem when using such a time scale. The intrinsic time  $\tau$  is only known for the past, contrary to the business time scale  $\vartheta$ , which is known also for the future, because it is based on average behavior. Thus a forecasting model for the price actually needs to be composed of *two* forecasting models, one for the intrinsic time and one for the price. The first requires forecasting of the size (not the direction) as time cannot flow backward.

<sup>5</sup> In the same paper Stock (1988) indicates how this approach can be compared to the ARCH models.

### 9.3.2 Model Structure

The price generating process is far from stationary in physical time. In Section 9.3.1, the geographic seasonality and conditional heteroskedasticity are modeled through successive changes of the underlying time scales. After these transformations, the remaining structure and the dynamics of the transformed series can be analyzed. The model presented in this section captures the dynamics through the computation of nonlinear indicators. Because the model is on the business-time scale  $\vartheta$ , all equations are written in terms of this new scale. The relation to the physical time scale is given by Equation 6.17.

### 9.3.3 A Linear Combination of Nonlinear Indicators

The model equations are based on nonlinear *indicators*, which are modeled with moving averages. Indicators for market prices come conceptually from simple trading systems used in practice by market participants.<sup>6</sup> Those 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, on the negative side as a *sell* signal. An indicator is thus used as a *predictor* of a variable or its change, for instance, a price change (i.e., a return).

Finding an ideal indicator, if it exists at all, would be enough to make a good price forecast. We, however, have no ideal indicators. Therefore there is need to combine different indicators appropriately to optimize their respective influence. Partly, the forecasting models presented here are based on a *linear combination* of price indicators  $z_x$  where the relative weights are estimated by multiple linear regression. For a fixed forecasting horizon  $\Delta\vartheta_f$  (corresponding to a  $\Delta t_f$  in physical time), the price forecast  $\tilde{x}_f$  is computed with

$$\tilde{x}_f = x_c + \sum_{j=1}^m c_{x,j}(\Delta\vartheta_f) z_{x,j}(\Delta\tilde{\tau}_f, \tau_c) \quad (9.8)$$

where  $x_c$  is the current price, and  $m$  is the number of indicators used in the model (from two to five per horizon). All the indicators are estimated in intrinsic time scale ( $\tau$ -scale). The coefficients  $c_{x,j}(\Delta\vartheta_f)$  are estimated with a multiple linear regression model.

$\Delta\tilde{\tau}_f$  in Equation 9.8, the forecasting horizon expressed in intrinsic time, is not yet defined. This quantity must be computed from its own forecasting model, which is similar to that in Equation 9.8. The forecasting horizon,  $\Delta\tilde{\tau}_f$ , can be written as an intrinsic time forecast,

$$\Delta\tilde{\tau}_f \equiv \tilde{\tau}_f - \tau_c = \sum_{j=1}^m c_{\tau,j}(\Delta\vartheta_f) z_{\tau,j}(\Delta\vartheta_f, \vartheta_c) \quad (9.9)$$

<sup>6</sup> See for instance Dunis and Feeny (1989); Murphy (1986).

where the forecasting model is computed in  $\vartheta$ -scale. The coefficients  $c_{\tau,j}(\Delta\vartheta_f)$  are estimated through a multiple linear regression and  $z_{\tau,j}(\Delta\vartheta_f, \vartheta_c)$  are the intrinsic time indicators.

Contrary to most traditional forecasting models, this model does not rely on a fixed basic time interval but is designed with a concept of continuous time. In fact, the time when a price is recorded in the database is unequally spaced in time. Moreover, the use of the  $\tau$ -scale implies that our forecasting models must be computed simultaneously over several fixed time horizons  $\Delta\vartheta_f$ .

Given a forecasting horizon in physical time  $\Delta t_f$  and the price history until  $x_c$ , one can compute  $\Delta\vartheta_f$  with Equation 6.17 and  $\tau_c$  with Equation 9.7. With a sufficiently large set of indicators,  $z_{\tau,j}(\Delta\vartheta_f, \vartheta_c)$ ,  $z_{x,j}(\Delta\tilde{\tau}_f, \tau_c)$ , and coefficients,  $c_{\tau,j}(\Delta\vartheta_f)$  and  $c_{x,j}(\Delta\vartheta_f)$ , the price forecast can be computed by choosing the appropriate  $\Delta\tilde{\tau}_f$  with Equation 9.9 and substituting it into Equation 9.8. In the next two sections, we define the indicators and study how to compute the coefficients  $c_{x,j}$ .

### 9.3.4 Moving Averages, Momenta, and Indicators

In Equations 9.8 and 9.9, the indicators are based on momenta, which are based on moving averages. In particular, we work with *exponential* moving averages (EMA) because they may be conveniently expressed in terms of recursion formulae (see Chapter 3).

The momentum indicator is


$$m_x(\Delta\tau_r, \tau_c) \equiv x_c - EMA[\Delta\tau_r; x](\tau_c) \quad (9.10)$$

which compares the most recent price to its own exponential moving average (EMA), using the notation in Section 3.3.5 and computed by the recursion formula of Equation 3.51. This is done by using intrinsic time as the time scale. It is also possible to define the *first* and the *second* momenta. The first momentum,  $m_x^{(1)}$ , is the difference of two exponential moving averages (or momenta) with *different ranges*. This can be considered as the first derivative of  $x(\tau)$ . The second momentum,  $m_x^{(2)}$ , is the linear combination of three exponential moving averages (or momenta) with different ranges, which provide information on the *curvature* of the series over a certain past history.

In section 9.3.3 we introduce the concept of indicators. Here we want to define those that are used in our forecasting models.

There are a large number of technical indicators (Murphy, 1986; Dunis and Feeny, 1989) and momenta indicators are widely used in technical trading systems. We limit our focus on momenta type indicators. Following Equation 9.8, let an indicator for returns with a range  $\Delta\tau_r$ ,  $z_x(\Delta\tau_r, \tau_c)$  be defined as follows:

$$z_x(\Delta\tau_r, \tau_c) = \left[ \frac{m_x^{(o)}(\Delta\tau_r, \tau_c)}{\sqrt{1 + (m_x^{(o)}(\Delta\tau_r, \tau_c)/m_{max})^2}} \right]^p \quad (9.11)$$



**FIGURE 9.2** The nonlinear function for computing the momentum indicator. This function is presented for different values of the parameter  $p$ .

where  $m_x^{(o)}(\Delta\tau_r, \tau_c)$  are normalized momenta of order  $o$  of returns,  $m_{max}^p$  is the maximum value the indicator can take, and the power  $p$  is the accentuator of the indicator movements. In the case of price indicators,  $p$  must be an odd number to keep the sign of the moving average. The shape of the nonlinear function is illustrated in Figure 9.2 for different powers  $p$  and for a  $m_{max}$  of one. This functional shape illustrates how the indicator plays the role of a primitive trading system. If the momentum has a high positive or negative value, the indicator  $z_x$  saturates, which is when the indicator is fully exposed in a *long* or *short* position. The power  $p$  both plays the role of a threshold (no threshold if  $p = 1$ ) and influences how the model approaches its full long or short position. The  $m_{max}^p$  value plays the role of the *quantity* of capital invested and also influences the shape of the indicator function.

The definition given in Equation 9.11 can easily be extended to other types of problems. For instance, the same definition can be used for constructing indicators for the intrinsic time in the  $\vartheta$ -scale,  $z_\tau(\Delta\vartheta_r, \vartheta_c)$  where the parameters are now defined as functions of  $\tau_c$  computed using Equation 9.7 on the  $\vartheta$ -scale. The function  $\tau(\vartheta)$  is a monotonic positive definite function so that all of its momenta are positive.<sup>7</sup> When the function is raised to an even power, only the upper right quadrant of Figure 9.2 becomes relevant. The primitive trading system analogy does not work in this case but the emphasis on large movements can be avoided by leveling off the indicator.

In the implementation of this algorithm, the indicators are *continuously updated*. Every new price received from the market makers causes the model to

<sup>7</sup> Time never flows backward.

recompute all its indicators for all the horizons. It then updates the forecasts for each time horizon.

### 9.3.5 Continuous Coefficient Update

The forecasting model environment is such that the indicators and the corresponding coefficients are continuously updated. Each coefficient  $(c_{x,j}, c_{\tau,j})$  is updated by estimating the model in the most recent past history. The length of the past history is a function of the forecasting horizon.<sup>8</sup> The motivation for horizon-dependent finite samples for optimization is motivated from the fact that there are different regimes in the market and short-term horizons are particularly sensitive to them. Furthermore, short-term traders are not influenced by a past much older than 3 months. The use of samples extending into the far past to optimize short forecasting horizons will make the model less adaptive to regime changes.

Adaptation to long-term regime or structural changes is enabled by re-evaluation of the optimization as soon as enough new information becomes available. The optimization sample size is kept fixed (in  $\vartheta$ -time), rolled forward, and then linear regression is reapplied to the new sample. This technique is similar to that used in Schinasi and Swamy (1989) and Swamy and Schinasi (1989) except that we use a fixed sample size while they add on the new data to their sample. The model is optimized through the usual generalized least-square method, except for two modifications.

Our forecasting models run in real time and the continuous reoptimization can generate instabilities (rapid jump from positive to negative forecast) when standard linear regression techniques are used. The instabilities originate from both the indicators and their coefficients in the linear combination (Equations 9.8 and 9.9). The indicators are moderately volatile and we avoid indicators that are too volatile by limiting the power of the exponents in the indicator construction to 3 for simple momenta and 7 for higher momenta. Moderately volatile indicators can cause instabilities only if their coefficients are large. The coefficients are less volatile than the indicators (due to the large optimization samples), but they may have high values if the regression by which they are optimized is near-singular because of high correlation between the indicators. Within a particular sample, the high positive and negative coefficients typical in the solution of a near-singular regression matrix would balance each other out. However, as soon as these coefficients are used with changing indicator values outside this sample, the equilibrium is lost and the high coefficients may boost the forecast signal. We have already eliminated one source of near-singularity by avoiding indicators that are too similar in the same forecast.

The standard regression technique is applied under the assumption of precise regressors and a dependent variable with a Gaussian error. Our regressors (indicators), however, originate from the same database as the dependent variable (the return); thus, they are prone to database errors (missing data, badly filtered data, and so on) and to errors in the construction of the  $\vartheta$  and  $\tau$  time scales. Taking into

<sup>8</sup> A few months for hourly forecasts, up to a few years for 3-month forecasts.

account the regressor errors allows a solution to the problem of near-singularities in a natural way. Instead of considering the  $j^{\text{th}}$  regressor  $z_{j,i}$  at the  $i^{\text{th}}$  observation (where we have dropped the variable index and the horizon for ease of notation), we consider the *imprecise* regressors  $\hat{z}_{j,i} = z_{j,i} + \varepsilon_{j,i}$  where  $\varepsilon_{j,i}$  is the random error with variance of  $\varrho^2$  times that of  $z_j$ . We call the small parameter  $\varrho$  the typical relative error of the indicators and we assume it is roughly the same for all indicators of the type we defined in Section 9.3.4. Without going into the details of the calculation, such a change modifies the final version of the system of  $k$  equations. The  $k^{\text{th}}$  equation can be written as follows:

$$\sum_{j=1}^m c_j (1 + \varrho^2 \delta_{jk}) \sum_{i=1}^N w_i \hat{z}_{j,i} \hat{z}_{k,i} = \sum_{i=1}^N w_i Y_i z_{k,i} \quad (9.12)$$

where  $N$  is the number of observations used in the regression,  $m$  is the number of indicators (same as in Equation 9.8),  $w_i$  is a weighting function depending on the type of moving averages used (here it is an exponential), and  $\delta_{jk}$  is the usual Kronecker symbol: 1 for  $j = k$  and 0 for  $j \neq k$ . The quantity  $Y_i$  is the usual response term of the regression:  $x(\vartheta_i + \Delta\vartheta_f) - x(\vartheta_i)$ . There is only one addition to the original regression: the diagonal elements of the system matrix are multiplied by a constant factor  $1 + \varrho^2$ , slightly greater than 1.

The effect of increasing diagonal values of the original matrix by the factor  $\varrho^2$  is to guarantee a minimum regularity of the modified matrix even if the original one is near-singular or even singular. The variable  $\varrho^2$  can be interpreted as the parameter of this minimum regularity. This desired effect is also accompanied by a slight decrease of the absolute values of the coefficients  $c_j$ , because the right-hand side of the equation system remains unaffected by the modification. The decreases are insignificant, the only exceptions being for coefficients inflated by near-singularity in the original regression: there, the absolute values decrease substantially, which is what we want anyway.

The other departure from the usual regression technique is a modification of the regression response  $Y_i$  necessitated by the leptokurtic behavior of returns. The forecast signals are much less leptokurtic than the returns, hence the optimization is dominated by exceptionally large real price movements rather than the “normal” price movements. This is also accentuated by the fact that it is squared returns that enter into the computation of the least square fit. Furthermore, the users of our forecasting models are more interested in the correct direction of the forecast than in the absolute size of a return forecast. A pure linear regression is thus inappropriate.

The minimization of the sum of squared deviations, however, has an important advantage: it can be reached by solving a system of linear equations. Theoretically though, the least sum of squares could be replaced by any utility function. Our problem is thus to find, within the framework of the regression technique, a more appropriate optimization (or utility) function. The best way to achieve this goal

FORECASTING RISK AND RETURN

is through a *mapping function* of the returns: the forecast should fit the mapped returns  $\hat{Y}_i$  rather than the *real* returns  $Y_i$ .

A suitable mapping function makes the mapped returns less leptokurtic than the original ones. The rest of the regression problem remains unchanged. The desired effects can be obtained with an underproportional mapping function presenting the following properties:

- Small returns should be amplified when considered in the regression, in order to establish a sufficient penalty against forecasts of the wrong direction.
- Large returns should be reduced when considered in the regression, so the distribution function of mapped returns is no longer leptokurtic.
- The mapping effects should decrease with the increasing time horizon size.

The choice of such a mapping function  $M$  is arbitrary provided it has the above properties. The one function used in our model is

$$\hat{Y}_i = M(Y_i) = \frac{AY_i}{[Y_i^2 + B]^\alpha} \quad (9.13)$$

with the parameters  $A$ ,  $B$ , and  $\alpha$  depending on the time horizon  $\Delta t_f$ . The same parameters are used for all different FX rates. They have been calibrated by trial and error in order to keep the full sample variance of the mapped returns on the same level as that of the original returns. The parameter  $\alpha$  must follow the condition  $0 \leq \alpha < 0.5$  because the mapping function must be an underproportional bijection.

#### 9.4 MEASURING FORECAST QUALITY

Two questions are relevant for testing any forecasting model of foreign exchange (FX) rates:

- What data should be used for testing?
- What is a good measure of forecasting accuracy?

Since the classical paper by Meese and Rogoff (1983), researchers in this field have been aware of the need for out-of-sample tests to truly check the forecast validity. Because of the statistical nature of FX rates, there would be little significance in the forecast accuracy measured on the same data that were used for optimizing the models. The real test comes when the model is run on data that were not used in constructing the model. In our case, our model being run in real-time, we have a continuous out-of-sample test. Besides the question of in and out-of-sample, there is a question as to what constitutes the relevant quantity for measuring the accuracy of forecasting methods. In Makridakis *et al.* (1983) the main measures are reviewed. We limit ourselves here to presenting the reasons as to why we chose certain types of measures and how we compute the uncertainty of these measures.



### 9.4.1 Appropriate Measures of Forecast Accuracy

Most standard measures rely on the mean square error (MSE) and the mean absolute error (MAE) for each time horizon. These errors are then compared to the similar ones produced by a naive forecasting model serving as a benchmark. One naive model may be the random walk forecast where expected returns are zero and the best forecast for the future is the current price. These accuracy measures are, however, all parametric in the sense that they rely on the desirable properties of means and variances, which occur when the underlying distributions are normal. The selection of the random walk model to derive the benchmark MSE or MAE is inherently inappropriate. It is in effect comparing the price volatility (MSE or MAE) with the forecasting error. There is no reason to expect that the heteroskedasticity and the leptokurticity of returns would not affect their MSE or MAE for a particular horizon. Thus the significance of their comparison with the forecast MSE or MAE is unclear. It might only reflect the properties of price volatility.

Such considerations have led us to formulate here nonparametric methods of analyzing forecast accuracy. These are generally “distribution free” measures in that they do not assume a normally distributed population and so can be used when this assumption is not valid. One measure that has this desirable property is the percentage of forecasts in the right direction. To a trader, for instance, it is more important to correctly forecast the direction (up or down) of any trend than its magnitude. We term this measure the *direction quality*, also known as the sign test:

$$D(\Delta t_f) \equiv \frac{N(\{\tilde{x}_f \mid (\tilde{x}_f - x_c)(x_f - x_c) > 0\})}{N(\{\tilde{x}_f \mid (\tilde{x}_f - x_c)(x_f - x_c) \neq 0\})} \quad (9.14)$$

where  $N$  is a function that gives the number of elements of a particular set of variables  $\{x\}$ , and  $x_c$  and  $\tilde{x}_f$  have the same definition as in Equation 9.8. We give the forecasting horizon in physical time  $\Delta t_f$  because the quality must be measured in the time scale in which people look at the forecasts. It should be clarified here that this definition does not test the cases where either the forecast is the same as the current price or when the price at time  $t_c + \Delta t_f$  is the same as the current one. To illustrate this problem, let us note that the random walk forecast cannot be measured by this definition. Other definitions could be used, like counting the case when the direction is zero as half right and half wrong. Excluding the cases where one of the two variables is zero would be a problem if this would occur very often. Our results show that it only occurs quite seldom and for the real signal  $(x_f - x_c)$  (few percent of the observations on the very short horizons) and almost never for the forecast signal  $(\tilde{x}_f - x_c)$ .

Unlike more conventional forecasts, for instance, a weather forecast, an FX rate forecast is valuable even when its direction quality is slightly above 50% and statistically significant. No trader expects to be right all the time. In practice we assume that a  $D$  significantly higher than 50% means that the forecasting model is better than the random walk. The problem lies in defining the word “significant.”

T O B S E R V E D " H I N T S



As much as one would like to be independent of the random walk assumption, we are still forced to go back to it in one way or the other, as here, when we want to define the significance level of the direction quality. As a first approximation, we define the significance level as the 95% confidence level of the random walk:

$$\sigma_D \approx \frac{1.96}{2\sqrt{n}} \quad (9.15)$$

where  $n$  is the number of tests. The factor 2 comes from the assumption of an equal probability of having positive or negative signals. It is a similar problem to the one of tossing a coin.

Another measure we use in conjunction with the previous one is the *signal correlation* between the forecasting signal and the real price signal:

$$C(\Delta t_f) \equiv \frac{\sum_{i=1}^{n'} (x_{f,i} - x_{c,i}) (\tilde{x}_{f,i} - x_{c,i})}{\sqrt{\sum_{i=1}^{n'} (x_{f,i} - x_{c,i})^2 \sum_{i=i'}^n (\tilde{x}_{f,i} - x_{c,i})^2}} \quad (9.16)$$

where  $n'$  is the number of possible measures in the full sample,  $n$  the number of full forecasting horizons in the full sample, and  $i'$  is  $n - n'$ . Here again the forecasting horizon is given in physical time,  $\Delta t_f$ . We estimate the significance of this quantity using  $1.96/\sqrt{n'}$ .

Both the direction quality and the signal correlation unfortunately have a slight drawback. They do not provide a measure of the forecast effectiveness. Nevertheless, we believe that they are superior to standard measures due to the nonnormality of the return distributions. The direction quality, which for all practical purposes is the most relevant indication of the forecast, and the signal correlation must be highly significant before we accept a model as being "satisfactory."

#### 9.4.2 Empirical Results for the Multi-Horizon Model

Optimization consists of two distinct, but interrelated operations, corresponding to the two main types of parameters in the models. The linear model coefficients  $c_{\tau,j}$  and  $c_{x,j}$  are optimized through least squares (see Section 9.3.5), and under the control of this process, always fulfill the strict out-of-sample condition when applied to a forecast. On the other hand, the nonlinear parameters of the indicators described in Section 9.3.4 must be optimized by trial and error to meet the above criteria: the direction quality and the signal correlation. The data set used in selecting the best combination of indicators is termed the *in-sample* period where the model parameters are fully optimized.

In Table 9.1 we indicate how our sample is divided to satisfy the different requirements of model initialization, in-sample optimization, and out-of-sample tests. The initialization period is needed for both initializing the different EMAs (see the discussion in Section 3.3.3) and computing the first set of linear coefficients  $c_{\tau,j}$  and  $c_{x,j}$ . The results presented in the next section are computed over two specific periods using our database of intraday market makers' quotes. The first

FOR RELEASE UNDER E.O. 14176

**TABLE 9.1** The sampling periods of the forecast study.

The data samples used for initialization, model training (in-sample), and testing (out-of-sample).

Data range	Data types	Data size	Usage
June 1, 1973, to Feb. 2, 1986	Daily data	152 months	Model initialization
Feb. 3, 1986, to Dec. 1, 1986	Intraday	10 months	Model initialization
Dec. 1, 1986, to Sep. 1, 1990	Intraday	46 months	In-sample period
Sep. 3, 1990, to Nov. 3, 1992	Intraday	25 months	Out-of-sample tests

runs from December 1, 1986, to September 1, 1990 (46 months), and is our in-sample period. The second runs from September 3, 1990, to October 3, 1992 (25 months). This second period is pure *post ex-ante* testing—that is, data from this period were not used at all for building the model. These 25 months constitute our out-of-sample test.

### 9.4.3 Forecast Effectiveness in Intraday Horizons

The forecast horizons here are for 2, 4, and 8 hr. This choice was made for two reasons. First, there is almost no literature to study models for such short horizons. Second, the statistical significance of the findings can be enhanced due to a large number of observations within a few years at the intraday frequency.

The quality measures are computed for each time horizon at an interval of 1/12 of the horizon. As mentioned earlier, the forecast accuracy is always measured in the physical time scale because it is in this scale that the different forecasts are useful. The number of relevant points in the statistical computation varies depending on the horizon and on the number of missing data for the different currencies. For the 2 hr horizon it varies from around 70,000 tests to 140,000, for 4 hr from 35,000 to 70,000, and for 8 hr 24,000 to 37,000. These very large numbers ensure that our statistical results are highly significant. We currently have 41 currencies running on the Olsen Information System (OIS), but in order not to overwhelm this study with numbers, we only show results for the 10 most important FX rates against the USD and 10 of the most traded cross rates. For the other currencies the results are very similar. The direction quality and the signal correlation are given in percentage for the in and out-of-sample testing periods in Table 9.2 for the USD rates and in Table 9.3 for the cross rates. Juxtaposing both results show clearly that, except for the 2 hr in the USD rates and for GBP-USD, the quality achieved in-sample is in most cases maintained out-of-sample and sometimes even slightly improved.

In Table 9.4 we summarize the significance of these results. For each horizon and for each currency we write a “+” sign if *both* the direction quality and the signal correlation are above the significance levels computed using Equation 9.15 for the direction quality and  $1.96/\sqrt{n'}$  for the signal correlation. If one of the two

FORBES "THE H200"

**TABLE 9.2** Direction quality and signal correlation for 10 USD rates.

Direction quality and signal correlation, in-sample and out-of-sample, for 9 FX rates and gold price against the USD. The numbers are expressed in percentage.

FX	Hor.	Direction	Correlation	FX	Hor.	Direction	Correlation
USD-DEM	2hr	52.1 / 51.6	+2.8 / +0.7	USD-NLG	2hr	51.5 / 50.8	+2.0 / -0.0
	4hr	52.6 / 52.1	+4.9 / +3.0		4hr	51.5 / 51.8	+1.9 / +1.2
	8hr	52.0 / 52.5	+2.8 / +2.0		8hr	50.5 / 51.8	+0.1 / +3.2
USD-JPY	2hr	51.5 / 52.2	+1.8 / +3.2	USD-ITL	2hr	51.5 / 50.8	+1.6 / -1.3
	4hr	51.7 / 52.9	+2.1 / +6.0		4hr	51.8 / 51.8	+2.4 / +1.4
	8hr	51.8 / 52.1	+4.0 / +5.4		8hr	51.7 / 52.4	+1.9 / +4.3
GBP-USD	2hr	51.8 / 50.5	+1.1 / -1.9	USD-CAD	2hr	51.8 / 52.6	+3.0 / +3.5
	4hr	51.5 / 51.7	+1.6 / +0.9		4hr	51.9 / 52.9	+4.8 / +5.2
	8hr	50.6 / 51.3	+1.9 / +2.1		8hr	52.4 / 53.6	+2.9 / +6.0
USD-CHF	2hr	51.5 / 51.2	+1.7 / -0.7	AUD-USD	2hr	51.8 / 53.4	+0.8 / +3.5
	4hr	52.2 / 52.1	+2.7 / +1.9		4hr	51.7 / 54.2	+2.7 / +6.5
	8hr	52.5 / 51.5	+2.4 / -0.4		8hr	51.4 / 53.7	+4.0 / +6.1
USD-FRF	2hr	51.0 / 50.8	+1.1 / -0.4	XAU-USD	2hr	52.5 / 53.0	+3.5 / +2.3
	4hr	52.2 / 51.6	+3.1 / +2.2		4hr	53.5 / 53.1	+4.3 / +3.8
	8hr	50.7 / 51.5	+0.1 / +0.4		8hr	53.5 / 52.9	+3.9 / +2.9

measures or both are below the significance level we write a “-” sign. Except for two USD rates (GBP-USD and USD-FRF), two cross rates (CAD-CHF and XAU-CHF), and the 2 hr horizon for the USD rates, which do not sustain conclusively the out-of-sample test, the other cases confirm the success of the model. The 2 hr cross-rates pass the out-of-sample test for 80% of the cases (only 40% of the cases for the USD rates) and the 4 hr for 80% and 8 hr for 90% of the cases. The USD rates for the 4 hr pass the test in 90% of the cases and for 8 hr in 80% of the cases.

In this chapter, we have shown that with the help of high-frequency data the statistical properties of FX rates can be better understood and that specifying forecasting models for very short-term horizons is possible. These models contain ingredients all designed to better capture the dynamics at work in the FX market. The most important characteristics of the models are as follows:

- Univariate time analysis type of model but based on intraday nonhomogeneous data,
- Variable time scales to capture both the seasonal heteroskedasticity ( $\vartheta$ -scale) and the autoregressive conditional heteroskedasticity ( $\tau$ -scale)
- Linear combination of nonlinear indicators
- Multiple linear regression with two modifications to avoid instabilities and to correct for the leptokurtic behavior of the returns

**TABLE 9.3** Direction quality and signal correlation for 10 cross rates.

Direction quality and signal correlation, in-sample and out-of-sample, for 9 FX cross rates and gold price. The numbers are expressed in percentage.

FX	Hor.	Direction	Correlation	FX	Hor.	Direction	Correlation
JPY-DEM	2hr	52.5 / 51.7	+2.3 / +1.0	GBP-CHF	2hr	54.4 / 55.1	+7.4 / +6.5
	4hr	51.8 / 51.0	+3.0 / +2.5		4hr	53.5 / 54.1	+6.2 / +5.0
	8hr	52.0 / 51.0	+4.0 / +1.3		8hr	53.8 / 54.2	+6.4 / +8.0
GBP-DEM	2hr	54.0 / 55.2	+6.1 / +5.6	JPY-CHF	2hr	52.6 / 51.7	+3.1 / +0.9
	4hr	53.3 / 54.5	+5.3 / +4.4		4hr	52.2 / 50.8	+3.5 / +2.1
	8hr	53.3 / 54.2	+4.7 / +8.7		8hr	54.1 / 51.4	+9.1 / +1.5
CHF-DEM	2hr	57.1 / 55.5	+12.5 / +8.1	GBP-JPY	2hr	52.8 / 53.3	+4.5 / +4.0
	4hr	54.9 / 54.6	+9.7 / +6.2		4hr	52.3 / 53.0	+4.6 / +4.8
	8hr	55.2 / 54.4	+10.0 / +6.7		8hr	52.8 / 52.8	+6.5 / +6.2
FRF-DEM	2hr	62.0 / 62.7	+20.4 / +22.0	CAD-CHF	2hr	51.1 / 51.2	+1.9 / -0.4
	4hr	59.7 / 59.5	+15.0 / +16.1		4hr	52.0 / 51.7	+2.6 / +0.4
	8hr	57.6 / 56.3	+14.3 / +13.2		8hr	52.3 / 51.8	+4.3 / +1.6
DEM-AUD	2hr	51.4 / 51.1	+2.5 / +1.2	XAU-CHF	2hr	51.0 / 50.6	+1.0 / -0.8
	4hr	51.2 / 51.1	+3.1 / +2.1		4hr	52.1 / 51.0	+3.2 / -0.2
	8hr	50.5 / 51.2	+4.0 / +3.2		8hr	52.0 / 52.4	+3.0 / +5.1

- Continuous optimization of the model coefficients in a finite size, forecasting horizon-dependent sample.

The forecast quality of these models is evaluated on a very large sample with two different measures that avoid statistical problems arising from the nature of the FX rate time series. The rigorous separation of in and out-of-sample measures, the large number of observations, and the stringent significance levels mean that the statistical results of the forecast evaluation are convincing evidence for our models having beat the random walk for most of the 20 studied currencies and for the very short-term forecasting horizons. These results are also corroborated by those we obtain on the other 21 rates that run on the Olsen Information System (OIS).

What are the consequences of such results on the economic theory of market efficiency? We believe that they point to the extension and improvement of methods and tools for defining and analyzing market efficiency. The accepted theory was probably never conceived for such short horizons, and even more important, it takes an unrealistic view of market response to new information. Being developed only in a statistical framework, the theory assumes that economic actors integrate new price information *instantaneously*, and very little attention is paid to the time needed for a piece of information to be available to all market participants and to the diverse interpretation of that information. In the context of very short time horizons these factors play critical roles in market adjustments. It is reasonable

**TABLE 9.4** Significance of the forecast quality for 20 exchange rates.

The in-sample and out-of-sample forecast significance for 10 USD rates and 10 cross rates. The "+" sign indicates a forecast quality above the significance limits of all test criteria, otherwise the "-" sign is used. Example: "+/-" means a significant in-sample quality and an insignificant out-of-sample quality.

2 hr	USD-DEM	+/-	USD-NLG	+/-	JPY-DEM	+/+	GBP-CHF	+/+
	USD-JPY	+/+	USD-ITL	+/-	GBP-DEM	+/+	JPY-CHF	+/+
	GBP-USD	+/-	USD-CAD	+/+	CHF-DEM	+/+	GBP-JPY	+/+
	USD-CHF	+/-	AUD-USD	+/+	FRF-DEM	+/+	CAD-CHF	+/-
	USD-FRF	+/-	XAU-USD	+/+	DEM-AUD	+/+	XAU-CHF	+/-
4 hr	USD-DEM	+/+	USD-NLG	+/+	JPY-DEM	+/+	GBP-CHF	+/+
	USD-JPY	+/+	USD-ITL	+/+	GBP-DEM	+/+	JPY-CHF	+/+
	GBP-USD	+/-	USD-CAD	+/+	CHF-DEM	+/+	GBP-JPY	+/+
	USD-CHF	+/+	AUD-USD	+/+	FRF-DEM	+/+	CAD-CHF	+/-
	USD-FRF	+/+	XAU-USD	+/+	DEM-AUD	+/+	XAU-CHF	+/-
8 hr	USD-DEM	+/+	USD-NLG	-/+	JPY-DEM	+/-	GBP-CHF	+/+
	USD-JPY	+/+	USD-ITL	+/+	GBP-DEM	+/+	JPY-CHF	+/+
	GBP-USD	+/+	USD-CAD	+/+	CHF-DEM	+/+	GBP-JPY	+/+
	USD-CHF	-/-	AUD-USD	+/+	FRF-DEM	+/+	CAD-CHF	+/+
	USD-FRF	-/-	XAU-USD	+/+	DEM-AUD	-/+	XAU-CHF	+/+

to assume that the markets need a *finite time* to adjust to any information and that this time depends on the nature of the information.

We think that these adjustments can be modeled and hence that a certain predictability of price movements exists. Our forecasting models, while a positive step in this direction, are nevertheless only a first one and there is still room for improvement through a better understanding and definition of intrinsic time and through the search for better indicators.

2025 RELEASE UNDER E.O. 14176

# 10

---

## CORRELATION AND MULTIVARIATE RISK

### 10.1 INTRODUCTION

Correlations and covariances between returns of different financial assets play an important role in fields such as risk management and portfolio allocation. This chapter addresses three problematic issues concerning linear correlation coefficients of returns, computed from high-frequency data:

1. The correlation of intraday, equally spaced time series derived from unevenly spaced tick-by-tick data deserves careful treatment if a bias resulting from the classical missing value problem is to be avoided. We propose a simple and easy-to-use method, which corrects for different data frequencies and gaps by updating the linear correlation coefficient calculation with the aid of covolatility weights. This is a bivariate alternative to time scale transformations which treat heteroskedasticity by expanding periods of higher volatility while contracting periods of lower volatility.
2. It is generally recognized that correlations between financial time series vary over time. We probe the stability of correlation as a function of time for 7 years of high-frequency foreign exchange rate, implied forward interest rate, and stock index data. Correlations as functions of time in

turn allow for estimations of the memory that correlations have for their past values.

3. It has been demonstrated that there is a dramatic decrease in correlation, as data frequency enters the intrahour level (the “Epps effect”<sup>1</sup>). We characterize the Epps effect for correlations between a number of financial time series and suggest its possible relation to tick frequency.

## 10.2 ESTIMATING THE DEPENDENCE OF FINANCIAL TIME SERIES

Measuring the dependence or independence of financial time series is of increasing interest to those concerned with multivariate decision formation (e.g., in risk assessment or portfolio allocation). Often this is estimated quantitatively using the linear correlation coefficient,<sup>2</sup> which is a basic measurement of the dependence between variables. Zumbach (1997) reviews many interesting measures of associations besides the linear correlation. The popularity of this measure stems from its simple definition, practical ease of use, and its straightforward results, which are easily interpreted, scale free, and directly comparable. Although the calculation of the correlation coefficient is well defined and rather simple, a number of unresolved issues exist with respect to application of the rule and interpretation of results in the high-frequency data domain.

- The data input for the correlation coefficient calculation are two time series with equal (i.e., homogeneous) spacing between ticks. This necessity is easily satisfied for low frequency ( $\leq$  one tick per week) data. However, the intraday case deserves more careful treatment if a resulting data bias is to be avoided. A problem arises when the two time series of unevenly spaced tick-by-tick data have different frequencies or active hours within a day, which may or may not overlap. We propose a simple and easy-to-use normalization method, which corrects for frequency differentials and data gaps. This alternative formulation updates the correlation calculation only where data exists, ensuring that there is no measurement bias resulting from the classical missing value problem (see Krzanowski and Marriott, 1994, 1995) or from differences in the active hours of the financial time series. In addition, this formulation remains scale free and straightforward to understand and implement.
- The linear correlation coefficient calculation largely discards the time variable. The variances of two time series and their covariance are constructed

<sup>1</sup> Epps (1979).

<sup>2</sup> The use of the linear correlation coefficient is appropriate not only for multivariate normal joint distributions but also for multivariate elliptical joint distributions. Many financial joint return distributions have been observed to fall into or close to this latter category. Also for fat-tailed return distributions, the linear correlation coefficient remains a useful and relevant measure of association; only the interpretation of results and, more specifically, the determination of accurate confidence limits is problematic. Correlations of *squared* returns from fat-tailed distributions are even more problematic.

FOR OFFICIAL USE ONLY

either with the assumption of being constant or as a type of average value if value changes are recognized. It is generally accepted that correlations in financial time series vary over time (Longin and Solnik, 1995) and are even subject to correlation “breakdown” or large changes in correlation in critical periods. In the discussion that follows, we probe the stability of correlation as a function of time, for a number of financial instruments, in order to determine the relevance of using high-frequency data. We go on to investigate the manner in which present correlation values are in turn correlated to their past values (autocorrelation of correlations). A model of the self-memory of correlation is proposed as the basis for the formulation of a long-term correlation forecast.

- The impact of time series data frequency on correlations should also be clearly established. This is especially relevant as higher frequency data becomes more widely available and more often used in order to improve statistics. Previous authors have demonstrated a dramatic decrease in correlation as data frequency enters the intra-hour level, for both stock (Epps, 1979) and foreign exchange returns (see Guillaume *et al.*, 1997; Low *et al.*, 1996). This discussion attempts to characterize and investigate more deeply the Epps effect in a number of financial time series through the examination of 7 years of high-frequency data.

### 10.3 COVOLATILITY WEIGHTING

The calculation of correlation coefficients is straightforward but some inconvenience is introduced via its simple definition. The correlation calculation requires two equally spaced (i.e., homogeneous) time series as input. This necessity is easily satisfied where low-frequency ( $\leq$  one tick per week) data are concerned. However, the problem requires more careful treatment at higher data frequencies where one cannot dictate the time or number of observations. One often faces two main problems when estimating correlation between two high-frequency time series. The first involves correlating two time series of inherently different frequencies. If the two time series are both regular with respect to data arrival intervals but of different frequencies, one might create from them two equally spaced, homogeneous time series, which both have frequencies equal to the lesser frequent of the two. This easy situation does not occur very often, though. It is more common to be faced with time series such as foreign exchange (FX) rates, where data frequency can vary from very few quotes to hundreds of quotes per hour. What is the best way to measure the dependence between an FX rate and another one that is perhaps less active or has activity peaks and valleys at completely different daytimes? Ideally, one would prefer the correlation calculation to be updated more often when more information exists and less often when it does not exist. A way to do this is to introduce a time scale that compresses physical time if there is no information and to expand it when it exists. This is similar to the idea presented in Chapter 6, where  $\vartheta$ -time was introduced to model volatility patterns. This method has been

FOR THE COURT



found useful for a number of applications, but is time-consuming to implement in practice. Moreover, we have the multivariate problem of two time series for which we would need a common time scale.

A second problem arising when estimating correlation between two high-frequency financial time series is that of missing values or data gaps. Large data gaps are actually an extreme case of the first problem (varying and nonmatching data arrival frequencies), but there is no harm in discussing the two problems separately. Despite one's best efforts, data gaps sometimes occur due to failures in the data acquisition chain. One can only make an educated guess about the correlation between two time series when such a gap occurs; it cannot be measured. More commonly, there are financial instruments whose time series have regular and large data gaps as part of their inherent character. Consider, for example, attempting to correlate a stock index (e.g., the Dow Jones Industrial Average,) which exists for 8 hr per day, 5 days per week (except holidays), with another stock index that exists for a similar amount of time each day but with a relatively large time shift (e.g., the Financial Times 100 index). There are a number of different schools of thought regarding the correlation between two financial instruments when one or both are not actually active. These sometimes consider derivatives of the instruments rather than the underlying instruments themselves. Other arguments confuse time-lagged correlation with direct correlation, but these are entirely different issues. When faced with varying activity rates and data gaps, it would be convenient to use some form of data interpolation to solve these problems. Unfortunately, the experience of many practitioners has not been reassuring (see Press *et al.*, 1992).

Some methods for approximating a homogeneous time series from unevenly spaced, tick-by-tick data involve some form of data imputation. Methods of imputing data vary in complexity and effectiveness and most have been found to be beneficial under at least some set of conditions and assumptions. However, all forms of imputation rely on a model, and a standard supposition is that critical characteristics of the data do not change between in-sample and out-of-sample periods. There is always the possibility that imputation will introduce a false bias into variance and covariance calculations, but nevertheless it is difficult to avoid some form of it in cases where data is not of an infinitely high frequency. Some useful attempts have been made to circumvent imputation all together. One interesting and recent example is described in de Jong and Nijman (1997). This work builds on efforts described in Cohen *et al.* (1983) and Lo and MacKinlay (1990a,b). The authors develop a covariance estimator, which uses irregularly spaced data whenever and wherever it exists in either of two time series. However, methods such as this one rest on the assumption that the processes generating transaction times and the prices themselves are independent. This assumption may be quite reasonable, depending on the instruments involved, but proving so is rarely trivial and we prefer to avoid it altogether.

In this discussion, we propose and illustrate a simple measure of correlation that avoids imputation based on data models or assumptions on distributional characteristics. Although the inputs for this alternative measure are homogeneous

FOR THE RECORD

time series derived through simple linear interpolation, the method filters out any underestimation of variances and covariances caused by lack of sampling variation. In addition, rather than making the strong assumption that price and transaction time are independent, this method makes use of the arrival time variable in order to compensate for the sometimes large differences that can exist in financial time series frequencies. Data gaps of varying size are common and we ignore any discussion of whether correlation actually exists during this period, because in any case we cannot measure it directly. Our goal is rather to develop a measure of correlation where information exists and to avoid updating our measure where data do not exist, a fact that should be recalled when results are interpreted. This implies that a lower data frequency or data gaps in one time series may limit the use of another one, and the unavoidable price to pay is a certain loss of statistical significance. However, the method is specifically meant to measure correlations at high data frequencies where statistical significance is high by nature.

### 10.3.1 Formulation of an Adjusted Correlation Measure

The standard linear correlation coefficient is a measure of correlation between two time series  $\Delta x_i$  and  $\Delta y_i$  and is defined as follows:

$$\rho(x_i, y_i) \equiv \frac{\sum_{i=1}^n (\Delta x_i - \langle \Delta x \rangle)(\Delta y_i - \langle \Delta y \rangle)}{\sqrt{\sum_{i=1}^n (\Delta x_i - \langle \Delta x \rangle)^2 \sum_{i=1}^n (\Delta y_i - \langle \Delta y \rangle)^2}} \quad (10.1)$$

with the sample means,

$$\langle \Delta x \rangle \equiv \sum_{i=1}^n \frac{\Delta x_i}{n} \quad \text{and} \quad \langle \Delta y \rangle \equiv \sum_{i=1}^n \frac{\Delta y_i}{n} \quad (10.2)$$

The sample is of size  $T$  with  $n = T/\Delta t$  homogeneously spaced observations. Correlation values are unitless and may range from  $-1$  (completely anticorrelated) to  $1$  (completely correlated). A value of zero indicates two uncorrelated series.

The two variables  $\Delta x_i$  and  $\Delta y_i$  are usually returns of two financial assets. In risk assessment (but not in portfolio allocation), the deviation of returns from the zero level is often considered instead of the deviation from the sample means  $\langle \Delta x \rangle$  and  $\langle \Delta y \rangle$ . In this special case, we can insert  $\langle \Delta x \rangle = \langle \Delta y \rangle = 0$  in Equation 10.1.

An estimate of the *local* covolatility for each of these observations is defined by further dividing each time span ( $\Delta t$ ) over which  $\Delta x_i$  and  $\Delta y_i$  are calculated into  $m$  equal subintervals from which subreturn values,  $\Delta \tilde{x}_j$  and  $\Delta \tilde{y}_j$ , can be obtained. This redefined time series now consists of  $\tilde{n} = T/\Delta \tilde{t}$  equally spaced return observations where  $\Delta t \equiv m\Delta \tilde{t}$ . The return definitions conform to Equation 3.7, based on logarithmic middle prices as in Equation 3.6. To obtain a homogeneous series, we need linear interpolation as introduced in Equation 3.2. The choice of linear interpolation method is essential.

For each of the previous coarse returns,  $\Delta x_i$  (as for  $\Delta y_i$ ), there exists a corresponding estimation of covolatility between the two homogeneous time series of

Downloaded from https://www.cambridge.org/core. University of Cambridge, on 05 Jun 2020 at 10:00:00, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9780511526352.010

returns

$$\omega_i(\Delta\tilde{x}_j; \Delta\tilde{y}_j; \Delta\tilde{t}) \equiv \sum_{j=1}^m (|\Delta\tilde{x}_{i,m-j} - \langle\Delta\tilde{x}_{i,m}\rangle| \cdot |\Delta\tilde{y}_{i,m-j} - \langle\Delta\tilde{y}_{i,m}\rangle|)^\alpha \quad (10.3)$$

where

$$\langle\Delta\tilde{x}_{i,m}\rangle = \sum_{j=1}^m \frac{\Delta\tilde{x}_{i,m-j}}{m} \quad \text{and} \quad \langle\Delta\tilde{y}_{i,m}\rangle = \sum_{j=1}^m \frac{\Delta\tilde{y}_{i,m-j}}{m} \quad (10.4)$$

The most obvious choice for  $\alpha$  is 0.5, though this can be investigated as a way to magnify or demagnify the weight given to farther outlying return values. A value of 0.5 is used in all cases described in this discussion.

Equation 10.3 formulates covolatility around the mean rather than around zero and it therefore follows that  $\omega_i = 0$  for the case of returns derived from two linearly interpolated prices existing outside of our region of interest,  $\Delta t$ . These covolatility estimates can be inserted as weights in all the sums computed to obtain the variances and covariance of the correlation calculation:

$$\tilde{q}(\Delta x_i, \Delta y_i, \omega_i) \equiv \frac{\sum_{i=1}^{T/\Delta t} [(\Delta x_i - \langle\Delta x\rangle)(\Delta y_i - \langle\Delta y\rangle)\omega_i]}{\sqrt{\sum_{i=1}^{T/\Delta t} [(\Delta x_i - \langle\Delta x\rangle)^2\omega_i]} \sqrt{\sum_{i=1}^{T/\Delta t} [(\Delta y_i - \langle\Delta y\rangle)^2\omega_i]}} \quad (10.5)$$

Note that  $\Delta x_i$  and  $\Delta y_i$  from Equation 10.5 are the same values as used in Equation 10.1, as they are logarithmic returns taken over the same time period,  $\Delta t$ . These coarse return values can then be defined as the sum of the fine return values

$$\Delta x_i \equiv \sum_{j=1}^m \Delta\tilde{x}_{i,m-j} \quad (10.6)$$

The sample means  $\langle\Delta x\rangle$  and  $\langle\Delta y\rangle$  have to be reconsidered in Equation 10.5. In the special case of risk assessment, we can still replace them by zero. Otherwise, we prefer that they are calculated again in a weighted fashion so that returns are considered only when observations over intervals of size  $\Delta t$  exist. Rather than keeping Equation 10.2, we define covolatility weighted mean values for both time series:

$$\langle\Delta x\rangle \equiv \frac{\sum_{i=1}^{T/\Delta t} (\Delta x_i \cdot \omega_i)}{\sum_{i=1}^{T/\Delta t} \omega_i} \quad \text{and} \quad \langle\Delta y\rangle \equiv \frac{\sum_{i=1}^{T/\Delta t} (\Delta y_i \cdot \omega_i)}{\sum_{i=1}^{T/\Delta t} \omega_i} \quad (10.7)$$

In this way, the means are calculated over the identically weighted data sample also used for the rest of the correlation calculation. The weights adjust for periods of lower or higher activity.

Equation 10.3 is formulated in such a way that  $\omega_i = 0$  for the case of returns interpolated over a data gap—that is, a tick interval that fully contains the analyzed interval of size  $\Delta t$ . Data gaps have no influence on the means, and the sums of Equations 10.5 and 10.7 are not updated there. The covolatility adjusted measure of correlation described by Equation 10.5 also retains the desirable characteristics of the original, standard linear correlation coefficient; it is scale free, and completely different measurements are directly comparable. In addition, this alternative method is only slightly more complicated to implement than the standard linear correlation coefficient and can easily be implemented on a computer.

As will be applied later, this correlation measure easily fits into the framework of autocorrelation analysis. Given a time series of correlations  $\tilde{q}_t$ , it can be correlated with a copy of itself but with different time lags ( $\tau$ ) between the two, as shown in Equation 10.8:

$$R(\tilde{q}(\Delta x_i, \Delta y_i, \omega_i), \tau) = \frac{\sum_{t=\tau+1}^n (\tilde{q}_t - \langle \tilde{q}_1 \rangle)(\tilde{q}_{t-\tau} - \langle \tilde{q}_2 \rangle)}{[\sum_{t=\tau+1}^n (\tilde{q}_t - \langle \tilde{q}_1 \rangle)^2 \sum_{t=\tau+1}^n (\tilde{q}_{t-\tau} - \langle \tilde{q}_2 \rangle)^2]^{1/2}} \quad (10.8)$$

for  $\tau > 0$ , where

$$\langle \tilde{q}_1 \rangle = \frac{1}{n-\tau} \sum_{t=\tau+1}^n \tilde{q}_t \quad \text{and} \quad \langle \tilde{q}_2 \rangle = \frac{1}{n-\tau} \sum_{t=\tau+1}^n \tilde{q}_{t-\tau} \quad (10.9)$$

For the discussions that follow, we measure correlation using the covolatility adjusted method described by Equation 10.5, unless otherwise stated, and always with  $m = 6$  and  $\alpha = 0.5$  (see Equation 10.3). Any subsequent use of the commonly recognized linear correlation coefficient (Equation 10.1) will be referred to as the “standard” method.

### 10.3.2 Monte Carlo and Empirical Tests

Various tests were performed on the covolatility adjusted correlation measure in order to test its behavior when applied to time series with differing frequencies and data gaps.

A first test used synthetic Monte Carlo data to illustrate the effectiveness of the method. Two separate, uncorrelated, normally distributed, i.i.d. random time series,  $A_i$  and  $B_i$ , were produced, each with zero mean, standard deviation  $\sigma = 0.01$  and size  $m = 10,000$ . A third series,  $C_i$ , can then be formed as a linear combination of the previous two:

$$C_{i=1}^m \equiv kA_{i=1}^m + (1-k)B_{i=1}^m \quad (10.10)$$

**TABLE 10.1** Results of a Monte Carlo simulation of correlations.

Comparing the covolatility adjusted linear correlation  $\tilde{\rho}$  to the standard linear correlation  $\rho$ , both applied to synthetic time series. The series  $D_i$  is like  $C_i$ , but regularly spaced sections of the data are replaced by linearly interpolated data. Details are described in the text. Note the similarity of the  $\rho(A_i, C_i)$  and  $\tilde{\rho}(A_i, D_i)$  columns.

Multiplier $k$ Equation 10.10	$\rho(A_i, C_i)$ Equation 10.1	$\rho(A_i, D_i)$ Equation 10.1	$\tilde{\rho}(A_i, D_i)$ Equation 10.5
0.0	0.00	0.00	0.00
0.1	0.12	0.10	0.12
0.2	0.23	0.15	0.22
0.3	0.38	0.28	0.38
0.4	0.52	0.40	0.51
0.5	0.69	0.51	0.69
0.6	0.83	0.62	0.82
0.7	0.92	0.67	0.91
0.8	0.97	0.72	0.95
0.9	0.99	0.74	0.97
1.0	1.00	0.74	0.99

where the constant  $k$  is selected such that  $0 \leq k \leq 1$ . In this way, the new series  $C_i$  has a controllable correlation to the original data series  $A_i$ .

The synthetic returns  $C_i$  were then cumulated to synthetic prices  $P_i$ , with starting value  $P_1 = 10$  and sample size  $m + 1 = 10,001$ :

$$P_{i=2}^{m+1} \equiv e^{\ln(P_{i-1}) + C_{i-1}} \quad (10.11)$$

The pure cumulation of  $C_i$  leads to synthetic logarithmic prices that are transformed to synthetic nonlogarithmic prices by the exponential function.

Repeated data sections, each consisting of 50 price observations, were then deleted in the time series  $P_i$  and replaced by prices linearly interpolated from the prices bracketing the deleted sections. The distance between these artificial data gaps also consisted of 50 observations, creating an alternating series of original data patches followed by data gaps filled with linearly interpolated prices. Finally, the first differences of this altered price series were taken to build a new series of returns,  $D_i$ .


Equation 10.5 was then used to measure the correlation between one of the original return distributions,  $A_i$ , and the manipulated return distribution,  $D_i$ , given various values of the constant multiplier  $k$ . Results are shown in comparison to the standard linear correlation calculation in Table 10.1.

A comparison of columns two ( $\rho(A, C)$ ) and four ( $\tilde{\rho}(A, D)$ ) shows that the covolatility adjusted correlation measure described by Equation 10.5 successfully approximates the original standard linear correlation between distributions  $A$  and  $C$  before some data patches were replaced by linearly interpolated values. Any

small deviations that exist are due to the statistical error ( $\sim 2\%$ ) of these tests. The third column of Table 10.1, by contrast, shows standard correlation values severely affected by the interpolation-filled data gaps. This simple example illustrates one of the original design goals of the covolatility adjusted linear correlation measure: correlation is measured where data exist, and the calculation is not updated where data do not exist.

Tests with foreign exchange data were performed to exemplify the effect of the covolatility adjusted correlation measure on time series with fluctuating data frequency and volatility. Homogeneous time series of USD-DEM prices were generated according to Section 3.2.2, equally spaced by 3-min intervals, once in physical time, once in  $\vartheta$ -time as explained in Chapter 6. USD-DEM has a high data frequency (see Table 2.2), but is also characterized by large intraday and intraweek fluctuations of both data frequency and volatility as shown in Section 5.6.2 and Figure 5.12. Absolute value of USD-DEM returns were used because they are known to have autocorrelations of greater magnitude than actual returns. Three autocorrelation functions are investigated: (1) standard autocorrelation (Equation 10.8) of 18-min returns in physical time, (2) standard autocorrelation of 18-min returns in  $\vartheta$ -time (see Chapter 6), and (3) autocorrelation measured by the covolatility adjusted correlation coefficients (Equation 10.5), analyzing 18-min in physical time. The covolatility computation was done in 3-min intervals and with  $m = 6$  (Equation 10.3), resulting in a covolatility value every 18 min. Results of these measurements are shown in Figure 10.1. A total data period of 6 months was used, ranging from January 1, 1996, to July 1, 1996.

The covolatility adjusted autocorrelation values (bullets in Figure 10.1) are significantly lower than the corresponding standard autocorrelation values, but close to the standard autocorrelation of the series equally spaced in  $\vartheta$ -time. We ascribe the high level of standard autocorrelation in physical time to the weekly seasonality of the data. The high absolute returns during working days and the low values on weekends are responsible for part of the high standard autocorrelation at lags up to about 1 day. As described in Chapter 6,  $\vartheta$ -time eliminates seasonality and thus the part of the autocorrelation caused by seasonality. The covolatility behaves similarly in the following respect. Weekends with their data gaps have extremely low covolatility values, so they are practically eliminated from the statistics. Weekly seasonality no longer affects the statistics. At lags around 24 hr, the picture is different. The covolatility adjusted autocorrelation approaches the value of the standard autocorrelation in a clear peak which indicates daily seasonality. Unlike  $\vartheta$ -time, which deforms time to eliminate seasonality, the covolatility adjusted correlation measure was designed to give a high weight to the most active periods, with no intention to hide all the seasonalities. Removing seasonality is not always desirable, so we find the covolatility adjusted correlation estimation to be a suitable method for many applications. In addition, the simplicity of this methodology lends itself to wider use.



**FIGURE 10.1** Autocorrelation of the absolute values of USD-DEM returns as a function of the time lag. The triangles ( $\Delta$ ) refer to standard autocorrelation of absolute returns, equally spaced in physical time. Bullets ( $\bullet$ ) refer to the covolatility adjusted autocorrelation of the same absolute returns. Crosses ( $+$ ) refer to standard autocorrelation of absolute returns, equally spaced in  $\vartheta$ -time. Sampling period: January 1, 1996, to July 1, 1996.

#### 10.4 STABILITY OF RETURN CORRELATIONS

When correlation is calculated between two time series, the assumption is that this quantity does not vary over time. For the case of financial time series this is seldom occurs, although time variance of the correlation coefficient over time can sometimes be small. This issue is critical for portfolio pricing and risk management where hedging techniques can become worthless when they are most needed, during periods known as correlation “breakdown,” or relatively rapid change. Boyer *et al.* (1997) have also demonstrated that a detection of correlation breakdown or other structural breaks by splitting a return distribution into a number of quantiles can yield misleading results. We use high-frequency data to estimate correlations literally as a function of time for a number of different financial time series in an effort to better understand the level of change that can occur. High-frequency correlation estimations are contrasted with lower-frequency estimates for the same sample periods. The “memory” that correlation coefficients have for their past values is also estimated for a number of examples using a simple

FOR THE "THESIS"



**TABLE 10.2** Data sampling for correlation as function of time.

Four different sampling schemes are selected to divide the total sampling period of size  $T$  from January 7, 1990, to January 5, 1997.

Correlation calculation period $T/N$	Data frequency (number of returns per day) $f = nN/T$	Number of returns per correlation calculation $n$	95% confidence band $1.96/\sqrt{n}$
365 days	1	365	0.10
128 days	3	384	0.10
32 days	12	384	0.10
7 days	72	504	0.09

and appropriate parameterization. Such estimations can be applied to long-term correlation forecasting, which is required, for example, to price or hedge financial options involving multiple assets, Gibson and Boyer (1997).

#### 10.4.1 Correlation Variations over Time

The general stability of correlation coefficients was examined using various correlation calculation intervals and data frequencies. This involved examination of a fixed historical time series over a time period  $T$ , from January 7, 1990, to January 5, 1997. The time series of returns ( $r(t_i)$ ) was then divided into  $N$  subsets of equal duration ( $T/N$ ) from which correlation coefficients were computed according to Equation 10.1. Four values of  $N$  were selected, while the total period  $T$  always remained constant. A homogeneous series of  $n$  returns was then chosen inside each period of size  $T/N$  via linear interpolation, so each correlation coefficient is based on  $n$  observations. Similar numbers  $n$  were selected for all the four values of  $N$  in order to maintain nearly uniform statistics, as shown in Table 10.2. In this table, the number of return observations per day,  $f = nN/T$ , is also given.

Results from these calculations are shown in Figures 10.2 to 10.7, where correlations versus time are displayed, and dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed random distributions. The confidence limits are slightly nonuniform due to small variations in statistics. Some correlations were computed with fewer observations than  $n$  because of missing observations. Whenever a weight  $\omega_i$  from Equation 10.3 was equal to zero, the corresponding observation was ignored. The weights  $\omega_i$  were not used for any other purpose, and the correlations remain standard linear correlations defined by Equation 10.1.

Correlation coefficient mean values and variances are given for each pair of financial instruments and for each of the four calculation frequencies in Table 10.3. Having virtually the same statistical significance for all correlation calculations shown in Figures 10.2 through 10.7, we can make a number of observations about

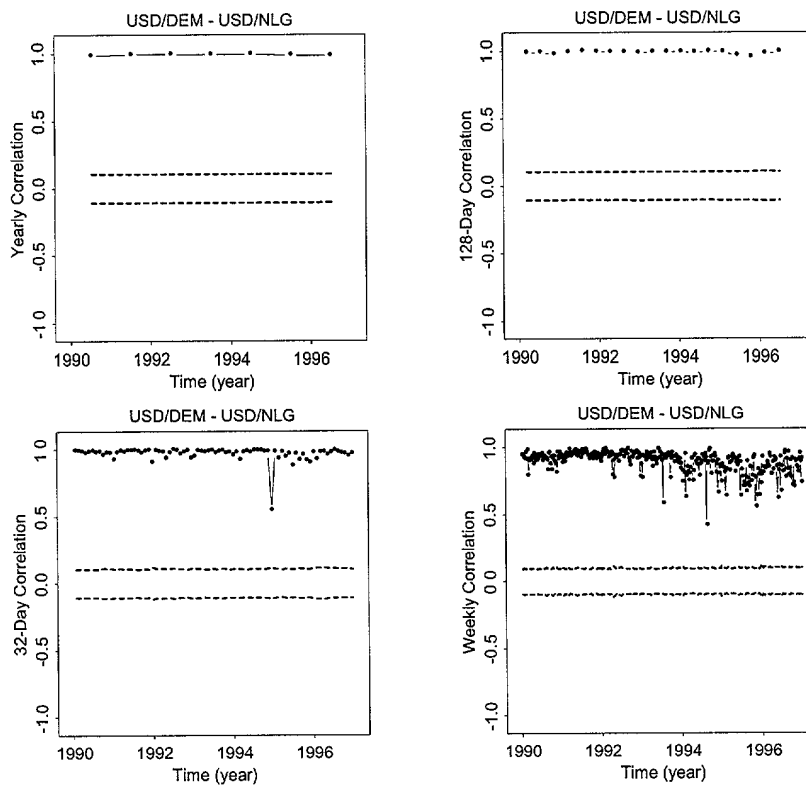


**TABLE 10.3** Means, variances, maxima and minima of correlation.

Means, variances, maxima and minima of the linear correlation coefficients as shown in Figures 10.2 through 10.7. For each pair of financial instruments, four correlation intervals  $T/N$  of decreasing size are investigated. The total sampling period  $T$  is from January 7, 1990, to January 5, 1997.

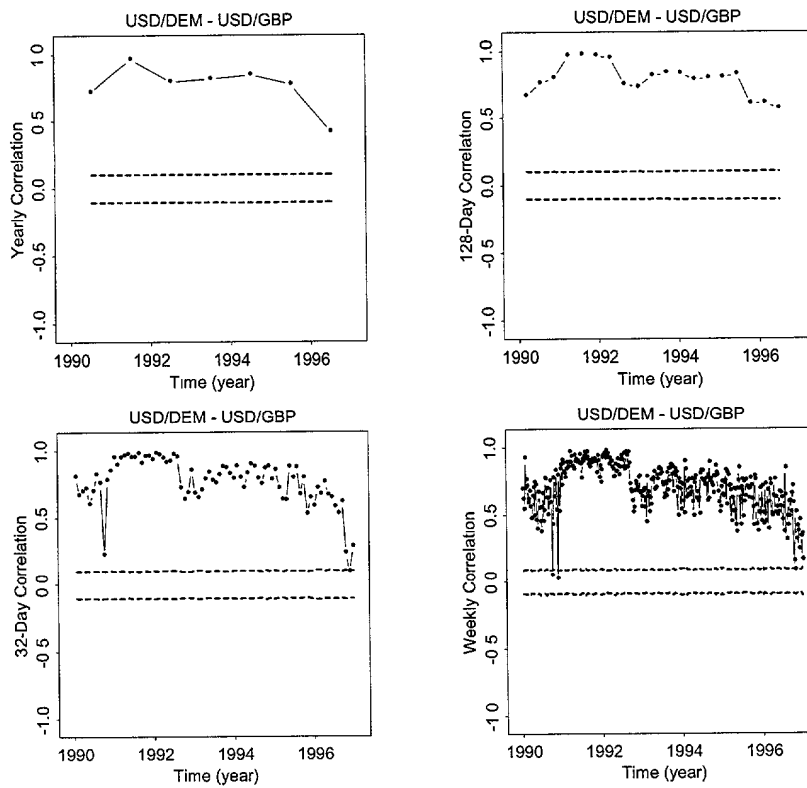
Instrument pair	Correlation period	Mean value	Variance ( $\sigma^2$ )	Max	Min.
USD-DEM – USD-NLG	1 year	0.99	0.000026	0.99	0.98
USD-DEM – USD-NLG	128 day	0.99	0.00012	1.00	0.95
USD-DEM – USD-NLG	32 day	0.96	0.0029	1.00	0.54
USD-DEM – USD-NLG	7 day	0.88	0.0067	0.98	0.41
USD-DEM – USD-GBP	1 year	0.76	0.029	0.96	0.42
USD-DEM – USD-GBP	128 day	0.79	0.015	0.98	0.57
USD-DEM – USD-GBP	32 day	0.76	0.031	0.98	0.09
USD-DEM – USD-GBP	7 day	0.69	0.030	0.97	0.20
USD-DEM – USD-ITL	1 year	0.76	0.040	0.99	0.41
USD-DEM – USD-ITL	128 day	0.75	0.057	0.99	0.18
USD-DEM – USD-ITL	32 day	0.76	0.044	0.99	0.07
USD-DEM – USD-ITL	7 day	0.68	0.044	0.97	0.07
DJIA – AMEX	1 year	0.73	0.0083	0.84	0.60
DJIA – AMEX	128 day	0.70	0.0087	0.85	0.57
DJIA – AMEX	32 day	0.41	0.041	0.78	-0.29
DJIA – AMEX	7 day	-0.01	0.030	0.62	-0.50
DEM 3-6m – DEM 9-12m	1 year	0.84	0.00074	0.88	0.81
DEM 3-6m – DEM 9-12m	128 day	0.78	0.0084	0.90	0.57
DEM 3-6m – DEM 9-12m	32 day	0.71	0.025	0.96	0.13
DEM 3-6m – DEM 9-12m	7 day	0.54	0.074	1.00	-1.00
USD 3-6m – DEM 3-6m	1 year	0.33	0.024	0.51	0.13
USD 3-6m – DEM 3-6m	128 day	0.30	0.028	0.59	-0.10
USD 3-6m – DEM 3-6m	32 day	0.30	0.051	0.85	-0.34
USD 3-6m – DEM 3-6m	7 day	0.28	0.066	1.00	-0.52

TABLE 10.3



**FIGURE 10.2** Linear correlation coefficients calculated using increasingly small sub-intervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the FX return pair USD-DEM – USD-NLG. The dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed returns.

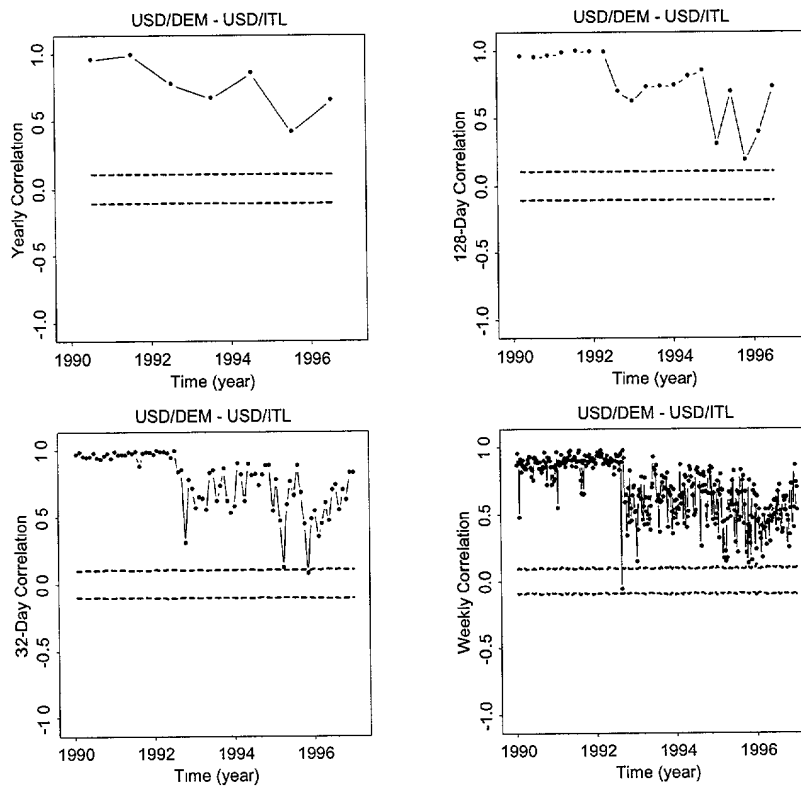
correlation stability. The highly correlated USD-DEM – USD-NLG returns shown in Figure 10.2 appear largely constant over the total sample period of 7 years. As the subperiod width for correlation calculation decreases (and the number of correlation calculations inside the total period increases), more structure becomes apparent. This additional structure is reflected by an increasing variance in Table 10.3. Correlations calculated with lower data frequency are not simply an average of those calculated with higher quotation frequencies; Table 10.3 shows the mean value for USD-DEM – USD-NLG correlations moving steadily downward with increasing correlation resolution (an -11% change between yearly data resolution and weekly resolution). This can be partially explained by considering that error distributions for empirically computed correlations are not symmetric when coefficients differ from zero. However, such drops in correlation with higher data frequency as can be observed with the DJIA-Amex pair point to a stronger effect which will be addressed in more detail in Section 10.5.



**FIGURE 10.3** Linear correlation coefficients calculated using increasingly small sub-intervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the FX return pair USD-DEM – USD-GBP. The dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed returns.

Figures 10.3 and 10.4 show correlations for the FX rate pairs USD-DEM – USD-GBP and USD-DEM – USD-ITL. Note that the inverse rate USD-GBP was used instead of the more usual GBP-USD, in order to be in line with other currencies and to obtain positive correlation. Both figures exhibit fast and large drops in correlations during the second and third weeks of September 1992. Presumably, this directly reflects the turmoil of the European Monetary System (EMS) at that time, when GBP and ITL left the system. This appears to be a clear example of correlation breakdown.

Recapitulating the major points of Section 10.4, correlation was examined as a function of time for a number of different financial instruments. The time frame of 7 years was divided in four different ways, using subperiods of 365, 128, 32, and 7 days. The number of subperiods over the total 7 year period was then 7, 19, 79, and 365, respectively. The subperiods of different size were divided into  $n$  intervals for return observations. This number  $n$  had roughly the same size in all cases,

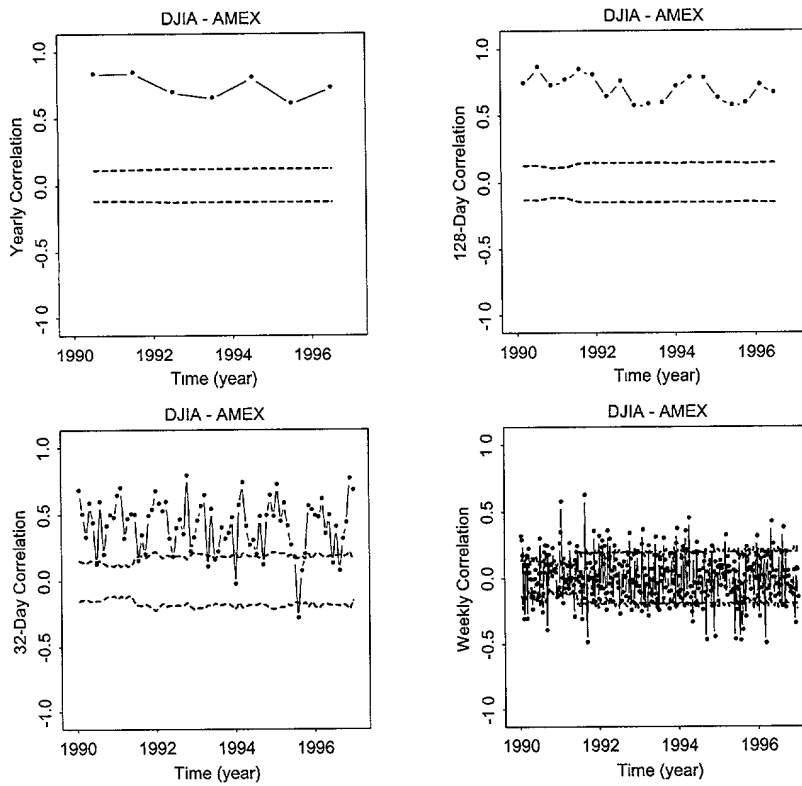


**FIGURE 10.4** Linear correlation coefficients calculated using increasingly small sub-intervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the FX return pair USD-DEM – USD-ITL. The dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed returns.

between 365 and 504, so the statistical significance of correlation values was always comparable. The correlation between some financial instruments can be described as reasonably stable. However, brief but large breaks can be observed in almost all cases, and the additional statistics provided by time series of higher frequency are essential to detect such occurrences. In addition, we observed decreasing absolute values all the correlations we examined when going to higher data frequencies. This will be further discussed further in Section 10.5.

#### 10.4.2 The Exponential Memory of Return Correlations

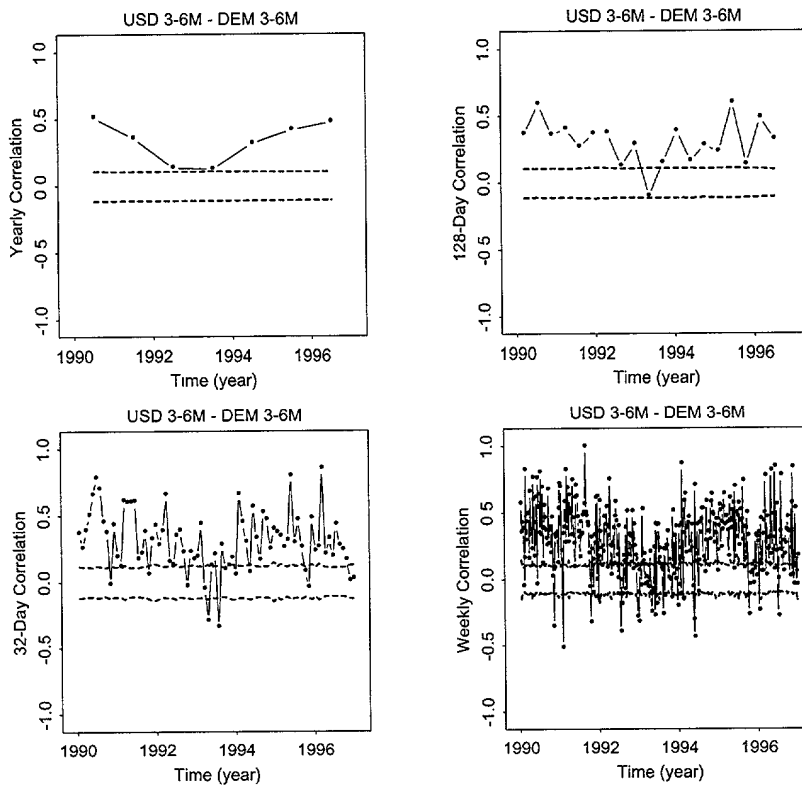
The linear correlation values shown in Figures 10.2 through 10.7 can be seen as time series in their own right. A certain stability of correlation levels seems to indicate that markets have a memory for these levels. This memory can be investigated by considering the autocorrelation of the correlation time series. We



**FIGURE 10.5** Linear correlation coefficients calculated using increasingly small sub-intervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the stock index pair DJIA – AMEX (both expressed in USD). The dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed returns.

focus on the weekly correlation measurements displayed in the lower right plots of Figures 10.2 through 10.7. These weekly correlations are computed from 20-min returns. The autocorrelation analysis was performed for different lags ( $\tau$ ), using Equation 10.8.

Results of these calculations are shown in Figure 10.8. Shown along with each autocorrelation curve are the 95% confidence limits for a normally distributed random process. The differences in the behaviors of the six correlation pairs are striking. For foreign exchange rate correlations, we observe a significantly positive autocorrelation extending to long lags up to 50 to 100 weeks. Correlation structures have a long memory. For correlations between implied forward interest rates, we find a positive autocorrelation above the significance limit for lags up to 3 or 4 months, which means a reduced but still long memory. The correlation of the stock index pair (Down Jones and AMEX Stock Index) behaves differently as it dives below significance already at the first lag of 1 week. The market has no consistent

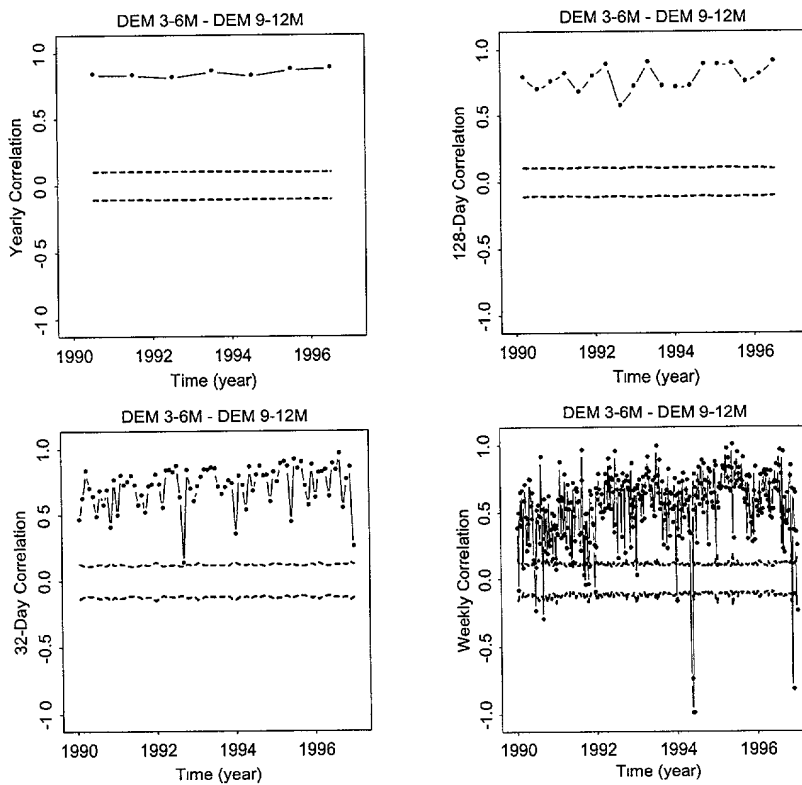


**FIGURE 10.6** Linear correlation coefficients calculated using increasingly small subintervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the implied forward interest rate pair USD 3-6 months – DEM 3-6 months. The dashed lines above and below zero correlation are 95% confidence intervals assuming Linear correlation coefficients calculated using increasingly small subintervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the implied forward interest rate pair USD 3-6 months – DEM 3-6 months. The dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed returns.

memory of the level, but this may be due to the strong Epps effect of the 20-min returns, explained in Section 10.5. Figure 10.8 also shows that autocorrelation values for each of the six instrument pairs decline roughly exponentially but with markedly different attenuation rates.

To better gauge the difference in autocorrelation attenuation for these correlation pairs, the autocorrelations shown in Figure 10.8 were modeled by a simple exponential function:


$$Y \equiv Ae^{-x/\lambda} \tag{10.12}$$



**FIGURE 10.7** Linear correlation coefficients calculated using increasingly small sub-intervals,  $T/N = (365 \text{ days}, 128 \text{ days}, 32 \text{ days}, \text{ and } 7 \text{ days})$ , for the implied forward interest rate pair DEM 3-6 months - DEM 9-12 months. The dashed lines above and below zero correlation are 95% confidence ranges assuming normally distributed returns.


where  $\lambda$  is an exponential attenuation length and  $A$  is a simple weight. These parameters were fitted to the data in Figure 10.8, starting with the first lag (neglecting the zeroth lag, which is equal to one by definition) and only to the point where autocorrelation data fell below the 68% of the upper confidence limit, thus focusing on the initial decay of autocorrelation. The autocorrelations of USD-DEM - USD-GBP correlations are shown in Figure 10.9. The results for the autocorrelations of correlation data shown in Figure 10.8 are also reported in Table 10.4. The DJIA-AMEX pair is excluded here due to the lack of correlation in the 20-minute returns, as mentioned above.

The goodness of fit can be judged by the  $\chi^2$  value, divided by the degree of freedom  $m$  in fitting. This is also shown in Table 10.4. Each autocorrelation value was assumed to have a stochastic error of  $1/\sqrt{N}$ , where  $N$  is the total number of correlation observations considered when calculating an individual autocorrelation value. All  $\chi^2/m$  values are below 1 or just slightly above 1, indicating



**FIGURE 10.8** Autocorrelations of weekly correlation coefficients of returns, as plotted in Figures 10.2 through 10.7. The 95% confidence ranges corresponding to normally distributed random distributions are shown as dotted curves, where the curvature is caused by the decreasing size of the sample with growing lags. The total sampling period  $T$  is from January 7, 1990, to January 5, 1997.





**FIGURE 10.9** Fit of an exponential function to the autocorrelation of USD-DEM – USD-GBP weekly correlation coefficients.

that Equation 10.12 describes the data rather well. Adding a second exponential function to Equation 10.12 did not significantly improve the goodness of fit in all cases unless the data point at zero lag (defined as being equal to one) was added to the data set.

Table 10.4 shows considerable values of the amplitude  $A$  (which cannot exceed 1) and a long memory of the correlation level. The pair USD-DEM – USD-NLG has the longest memory with an exponential attenuation length of more than 80 weeks, which is just over one and a half years. The autocorrelation model of Equation 10.12 with the parameters of Table 10.4 can be the basis of a correlation forecast. This forecast could be remarkably long-term due to the long memory in correlation, depending on the instrument pair.

### 10.5 CORRELATION BEHAVIOR AT HIGH DATA FREQUENCIES

Previous authors have observed a dramatic decrease in correlation as the time intervals of the returns enter the intrahour level, for both stock (Epps, 1979) and FX returns (Guillaume *et al.*, 1997; Low *et al.*, 1996). We follow the suggestion of Low *et al.* (1996) by referring to this phenomenon as the Epps (1979) effect after the first identifiable author to thoroughly document it. In this discussion, the Epps

**TABLE 10.4** Autocorrelation study.

The autocorrelations of weekly correlation data shown in Figure 10.8 are fitted to the parametrization given in Equation 10.12, for five financial instrument pairs. A large value of  $\lambda$  indicates a long, exponentially decaying memory of the correlation level. The variable  $\chi^2/m$  indicates the goodness of fit (good fits have a value around 1 or preferably below);  $m$  is the degree of freedom of fitting.

Instrument	$A$	$\lambda$ (weeks)	$\chi^2/m$
USD-DEM – USD-NLG	$0.35 \pm 0.01$	$80.9 \pm 7.1$	0.65
USD-DEM – USD-GBP	$0.62 \pm 0.02$	$31.5 \pm 1.8$	0.15
USD-DEM – USD-ITL	$0.61 \pm 0.02$	$59.8 \pm 3.0$	0.34
DEM 3-6m – DEM 9-12m	$0.27 \pm 0.04$	$10.0 \pm 2.6$	1.30
USD 3-6m – DEM 3-6m	$0.23 \pm 0.03$	$21.8 \pm 5.4$	0.41

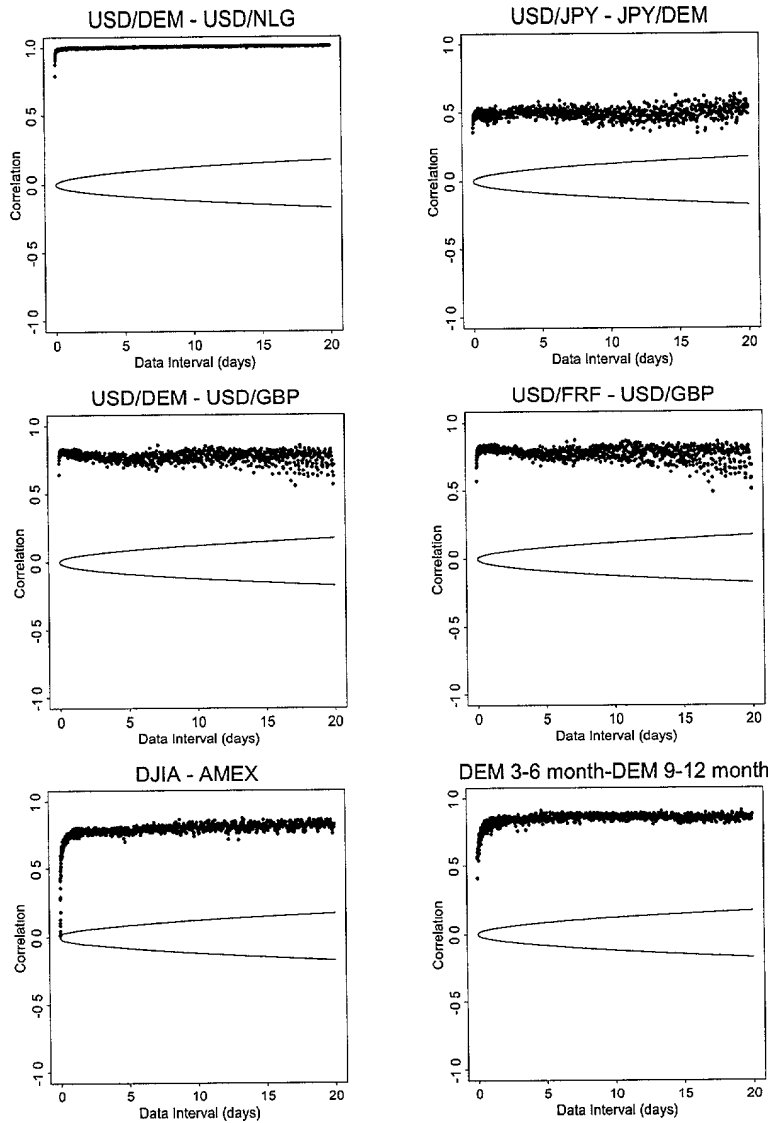
effect is characterized and investigated for a number of foreign exchange rates, stock indices, and implied forward interest rate pairs through the examination of the same 7 years of high-frequency return values as used in the previous sections.

The basis of our exploration was a set of homogeneous time series of returns, equally spaced by 5-min intervals, for several financial instruments. Linear interpolation was used in the sense of Equations 3.2 and 3.6. For each time series of 5-min returns, we obtained 499 additional time series through aggregation: 10-min returns, 15-min returns, . . . , 2500-min returns. To cover longer time intervals, 877 more time series were obtained by further aggregation in coarser steps: 2530-min returns, 2560-min returns, . . . , 28810-min returns, where a 28810-min interval roughly corresponds to 20 days. Various calculations were performed with these many times series. The most interesting results are the correlation coefficients for returns of different financial instruments, using the same time interval size. This can be done for all interval sizes of the aggregated time series.

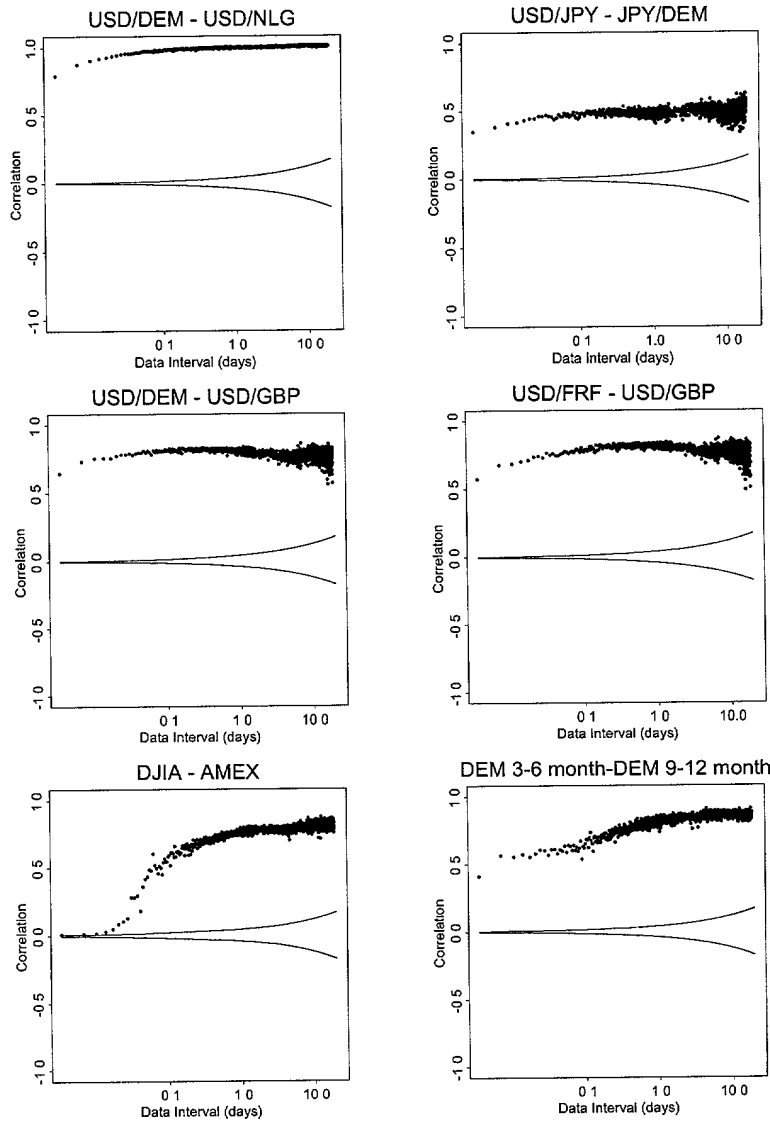
The resulting correlation coefficients are plotted as a function of the time interval size in Figure 10.10. When returns are computed with high frequency, over intervals distinctly shorter than 1 day, the correlation levels diminish in Figure 10.10. The same effect is better viewed in Figure 10.11, where the same data are shown with a logarithmic scale of interval sizes and where the data point farthest to the left (highest data frequency) corresponds to the correlation calculated using linearly interpolated, homogeneous time series of 5-min returns.

Table 10.5 gives the minimum and maximum values for the linear correlation coefficient data shown in Figure 10.10. Also given are the time intervals at which maxima occurred. We noted several problems when taking the maximum of correlation as a function of the time interval. The correlation graphs reach a more or less stable maximum level at time intervals exceeding one day, but this stability is not perfect for any correlation graph. Moreover, the maximum of correlation is affected by increasing stochastic noise for large time intervals. In an attempt

FOR THE COURT REPORTER



**FIGURE 10.10** Linear correlation coefficients calculated for six pairs of financial instruments as a function of the size of the time interval of returns. For all calculations the total sampling period remained constant (from January 9, 1990, to January 7, 1997,) causing the 95% confidence ranges to be narrow at high data frequencies and wider as the time interval increases. Rapid declines in correlation at very high data frequencies are noted in all cases.



**FIGURE 10.11** Linear correlation coefficients calculated for six pairs of financial instruments as a function of the size of the time interval of returns. The same data as in Figure 10.10 are shown with a logarithmic horizontal axis. Rapid declines in correlation at very high, intraday data frequencies are noted in all cases.


**TABLE 10.5** Correlation results characterizing the Epps effect.

Minima, maxima, and mean values (averaged over time intervals between 1 and 2 days) of correlation coefficients. These correlations are functions of the time interval of return measurement, as plotted in Figures 10.10 and 10.11. Also given are the time intervals at which the maxima occurred and the time intervals where the coefficient reaches 90% of the mean value. This latter time interval in the last column is called the stabilization interval—the threshold after which the Epps effect no longer affects correlation results. The sampling period was from January 9, 1990, to January 7, 1997.

Instrument pair	Min. corr.	Max. corr.	Interval of max. (days)	Mean corr. (1–2 days)	90% of mean	Stabiliz. interval (min)
USD-DEM – USD-GBP	0.55	0.86	7.2	0.79	0.71	10
USD-DEM – USD-NLG	0.78	1.00	14.0	0.99	0.89	15
USD-FRF – USD-ITL	0.49	0.86	12.0	0.79	0.71	25
USD-NLG – USD-FRF	0.69	0.99	16.9	0.97	0.87	25
USD-FRF – USD-GBP	0.48	0.86	7.2	0.80	0.72	30
USD-JPY – DEM-JPY	0.34	0.62	19.4	0.48	0.43	30
DEM-GBP – USD-GBP	0.23	0.75	17.0	0.45	0.41	170
DJIA – AMEX	0.00	0.86	13.3	0.77	0.69	320
DEM 3-6M – DEM 9-12M	0.40	0.90	19.2	0.82	0.74	340

to give a more accurate reference level of the correlation drop due to the Epps effect, Table 10.5 also reports the arithmetic mean of all correlation values based on time intervals between 1 and 2 days. A total of 224 correlation values (i.e., 224 aggregated time series) belong to this range of time intervals. Although there is no best choice of time interval for the general case, this mean value is considered as a reasonable reference level of correlation in its stable region. When moving to shorter time intervals, the Epps effect makes the correlation values decline. As a threshold value of the Epps effect, Table 10.5 also shows the time interval at which correlations drop to 90% of the reference level. This estimation of the Epps threshold or, seen from the other side, the correlation stabilization interval has the advantage that it can be uniformly applied to all cases, and it does not deliver obviously misleading results based on the maximum value of correlation.

We find that even currency pairs that are highly correlated in the long term become much less correlated in the intrahour data frequency range. Müller *et al.* (1997a) propose a hypothesis of heterogeneous markets where the market agents differ in their perception of the market, have differing risk profiles, and operate under different institutional constraints (see also Section 7.4). If the financial markets are indeed composed of heterogeneous agents with different time horizons of interest, then the Epps effect in correlation estimations can be interpreted as a cut-off between groups of agents. Short-term traders focus on the rapid movements of individual rates rather than multivariate sets of assets. For these short-term agents, correlations between instruments play a secondary role. Other, less rapid agents reestablish “correct” correlations after market shocks, but this takes some time.



**FIGURE 10.12** Correlation stabilization intervals (= stabilization points, in min) as a function of the inverse square root of the product of the tick frequencies. The same instrument pairs and stabilization intervals as in Table 10.5 are plotted.

When considering the stabilization interval in the last column of Table 10.5, we find that the Epps effect already vanishes at return measurement intervals of 10 min for the correlation of the most frequently quoted financial instruments. For pairs of less frequently quoted instruments, the Epps effect may last for hours, up to 6 hr in Table 10.5. This indicates a relationship between tick frequencies (and perhaps liquidity) of instruments and the duration of the Epps effect. Two correlation studies were made to probe this relationship: (1) standard correlation between the stabilization interval and the tick frequency of the more frequently quoted instrument of the pair and (2) standard correlation between the stabilization interval and the tick frequency of the less frequently quoted instrument. Mean tick frequencies per business day were taken, as listed in Table 2.2.

The greater of the two tick frequencies is estimated to have a standard correlation of -0.59 to the stabilization interval. The corresponding standard correlation of the lower tick frequency is -0.65. These values are significant to 92% and 95% confidence levels, respectively, assuming a normal random distribution. Therefore we conclude that, to a reasonable level of confidence, both of the tick frequencies substantially affect the stabilization interval after which the Epps effect of correlations vanishes. Tick frequency and stabilization interval are inversely related. This can be seen graphically in Figure 10.12 where the stabilization interval is plotted versus the inverse geometric mean of the two tick frequencies. When two tick frequencies are very high, as on the left-hand side of Figure 10.12, the stabilization interval becomes small. On the other side, at very low data frequencies, a plateau in the correlation stabilization interval appears to exist at a data interval of 300 to 400 min. This would indicate that the Epps effect does not play a substantial role in attenuating correlation values beyond 6-hr return measurement intervals,

FORBIDDEN TO REPRODUCE

even if the instruments involved are very inactive ( $< 100$  data updates per business day). It should be stated that these are indicative and preliminary results and an enhanced study with more instruments and statistics is called for.

## 10.6 CONCLUSIONS

The problems associated with estimation of correlation at higher data frequencies have been discussed and illustrated using examples. An easy-to-use covolatility adjusted correlation estimator, which correctly accounts for missing or nonexistent data, has been proposed. The effect of this new formulation is to estimate correlation when data exist, and to make no update to the correlation calculation when data do not exist. The input of the method is homogeneous time series linearly interpolated between the ticks. At times when tick intervals are longer than the return measurement intervals, the weight of the return observations tends to zero. Because the estimator is adjusted by covolatility, some of the information from the more frequent of the two time series involved will not be fully utilized, and statistical significance can be degraded. With growing data frequency, this degradation is inevitable but tolerable because the statistical significance based on high-frequency sampling is high by nature. Covolatility adjusted correlation is an estimator complementary to other estimators. Its fluctuating weighting of observations is an alternative to time scale transformations such as the  $\nu$ -time discussed in Chapter 6.

Empirically estimated linear correlation coefficients of returns vary over time. The return correlations of some financial instrument pairs widely fluctuate from week to week, whereas other correlations are very stable over periods of many years. It was observed that long-term historical stability is not a guarantee of future correlation stability. This was evidenced through the examination of USD-ITL and USD-GBP return correlations with the USD-DEM rate. The crisis of 1992, when the involvement of ITL and GBP in the European Monetary System (EMS) was suspended, was reflected by a dramatic and rapid change in their correlations with other currencies thereafter. Correlation calculated over a long data sample (years) has an averaging effect, and increased structure in correlations is observed when correlations are calculated over smaller periods (weeks). Depending on the time horizon of interest, there is pertinent information to be gained when calculations are performed over smaller periods using high-frequency data. The self-memories of return correlations have been modeled as exponential attenuations through estimation of the autocorrelation of linear correlation coefficients. Correlation values have memories for their past values that differ between instrument pairs and often extend over years rather than only weeks. The understanding of this correlation memory is a first step toward correlation forecasting.

The behavior of the correlation coefficient as a function of the time interval of return measurement has been investigated. Nonzero correlations of returns are dramatically attenuated when this interval decreases and enters the intrahour region. This behavior is called the Epps effect and depends on the pair of investi-

2025 RELEASE UNDER E.O. 14176

gated financial instruments. When the measurement interval exceeds a threshold value called the stabilization interval, the Epps effect gives way to a rather stable behavior of the correlation. There is some preliminary evidence of an inverse relationship between the stabilization interval and the mean tick frequency of the instruments involved. If financial markets are composed of heterogeneous agents as suggested in Section 7.4, the stabilization interval can be interpreted as a threshold between groups of agents. For extremely short-term traders focusing on time horizons below the stabilization interval, correlations between instruments may be less of an issue than for other agents.

In applications such as asset allocation or risk assessment, the return measurement intervals should preferably be chosen longer than the stabilization interval. However, there is no general “best” time interval for measuring correlation. It is important to choose the most relevant interval size for a specific application.

0393



# 11

---

## TRADING MODELS

### 11.1 INTRODUCTION

Recently, the skepticism among academics to the possibility of developing profitable trading models has decreased with the publication of many papers that document profitable trading strategies in financial markets, even when including transaction costs.

In the earlier literature, simple technical indicators for the securities market have been tested by Brock *et al.* (1992). Their study indicates that patterns uncovered by technical rules cannot be explained by simple linear processes or by changes in the behavior of volatility.<sup>1</sup> LeBaron (1992a), LeBaron (1997) and Levich and Thomas (1993b) follow the methodology of Brock *et al.* (1992) and use bootstrap simulations to demonstrate the statistical significance of the technical

---

<sup>1</sup> In Gençay (1998b), the DJIA data set of Brock *et al.* (1992) is studied with simple moving average indicators within the nonparametric conditional mean models. The results indicate that nonparametric models with buy-sell signals of the moving average models provide more accurate sign and mean squared prediction errors (MSPE) relative to random walk and GARCH models. Gençay (1999) shows that past buy-sell signals of simple moving average rules provide statistically significant sign predictions for modeling the conditional mean of the returns for the foreign exchange rates. The results in Gençay (1999) also indicate that past buy-sell signals of the simple moving average rules are more powerful for modeling the conditional mean dynamics in the nonparametric models.

trading rules against well-known parametric null models of exchange rates. Sullivan *et al.* (1999) examine the trading rule performance by extending the Brock *et al.* (1992) data for the period of 1987–1996. They show that the trading rule performance remains superior for the time period that Brock *et al.* (1992) studied; however, these gains disappear in the last 10 years of the Dow Jones Industrial Average (DJIA) series. Lo *et al.* (2000) have proposed an approach to evaluate the efficacy of technical analysis based on technical pattern recognition using nonparametric kernel regression. They apply their method to a large number of U.S. stocks and they report that several technical indicators provide incremental information of practical value. Overall, the scope of the most recent literature supports the technical analysis, but it is generally limited to simple univariate technical rules. One particular exception is the study by Dacorogna *et al.* (1995), which examines real-time trading models of foreign exchanges under heterogeneous trading strategies. They conclude that it is the identification of the heterogeneous market microstructure in a trading model which leads to an excess return after adjusting for market risk.

Trading models are investment tools that provide explicit buy and sell trading recommendations. A clear distinction should be made between a price change forecast (presented in Chapter 9) and an actual trading recommendation. A trading recommendation naturally includes a type of price change forecast, but must also account for the specific risk profile of the dealer or user of the respective trading model. Another distinction is that a trading model must take into account its past trading history. This decision might be biased by the position it is currently holding and the price paid for entering in this position, whereas a price forecast is not submitted to such asymmetries. A trading model thus goes beyond predicting a return. It must decide if a certain action is to be taken. This decision is subject to the specific risk profile, the trading history, and institutional constraints such as opening hours or business holidays.

The purpose of this chapter is not to provide *ready-to-use* trading strategies, but to give a description of the main ingredients needed in order for any real-time trading model to be usable for actual trading on financial markets. Any reasonable trading strategy is composed of a set of tools that provides trading recommendations within a capital management system. In this book we shall not discuss the capital management part, but we wish to show that with a reasoned approach and high-quality data, it is possible to design practical and profitable trading models. Indeed, we have developed our own trading models and this presentation builds on this experience. Our models anticipate price movements in the foreign exchange (FX) market sufficiently well to be profitable for many years yet with acceptable risk behavior, and, they have been used by many banks.

Market investors mainly use trading models as decision tools, but in this chapter we will also illustrate that profitable trading models with robust performance measures can be employed as a statistical tool to study the market structure and to test the adequacy of price-generation processes.

A robust performance measure of trading strategies is one of the most important ingredients in the development of new models and also in their use. In Section 11.3, we discuss different possible performance measures and we derive two risk-adjusted ones for investors with risk-averse preferences. Maximizing these measures is equivalent to maximizing the expected utility of an investor.

To construct successful trading strategies is not an easy task and many possible mistakes must be avoided during the different development phases of new models. We shall describe here some of the main traps in which new system designers generally fall and provide some ideas as to how to construct more robust trading strategies. In the following sections we will also give a short description of the various components needed in trading models and a specific approach using genetic algorithms to obtain more robust optimization results.

## 11.2 REAL-TIME TRADING STRATEGIES

In the assumption of a heterogeneous market, there is no trading strategy that is absolutely better than other ones. Which strategy to choose will depend on the trading and risk profile of the investor. This is confirmed by the existence of many different types of portfolio and investment strategies in the financial markets. It is also why we use in this study different trading model algorithms. We believe that these new investment strategies will simply contribute to, and not fundamentally change, the heterogeneous composition of financial markets.

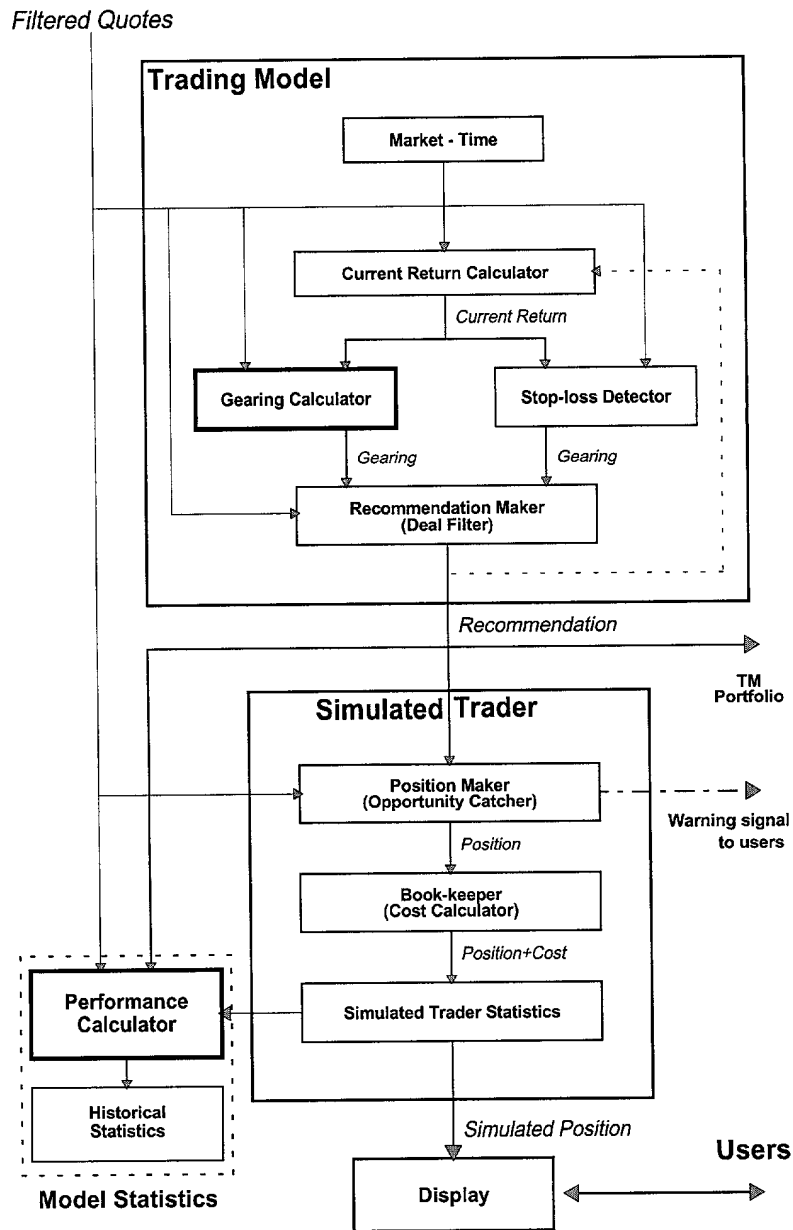
To be useful, real-time trading models must provide realistic trading recommendations that the user can follow. This means that the models should do the following:

- Give a warning a few minutes in advance of a deal.
- Not change recommendations too rapidly.
- Not give recommendations outside business hours.
- Take into account market holidays.
- Support stop-loss (around the clock).

In this section, we present the basic system architecture that we use in our real-time trading models and discuss the main components needed to transform available price quotes into actual trading recommendations. The model is divided in three main parts, that is,

- Generation of the *trading model* recommendations.
- Receipt of the simulated positions by the *simulated trader*.
- Generation of the model statistics by the *performance calculator*.

Figure 11.1 depicts the overall structure and data flow of a simple real-time trading model. The next subsections describe these different components.



**FIGURE 11.1** Data flow of prices and deal recommendations within a real-time trading model.

**TABLE 11.1** Market constraints.

FX market business time constraints for the trading models running for different geographical markets. The markets are listed roughly in the order of their opening times in GMT.

Market	Time zone	Opening time (local time)	Closing time (local time)	Holidays (per year)
Tokyo	JPT	09:00:00	18:00:00	15 Days
Singapore	PRC	09:00:00	18:00:00	11 Days
Frankfurt	MET	08:30:00	17:00:00	12 Days
Vienna	MET	08:00:00	17:30:00	15 Days
Zurich	MET	08:00:00	17:30:00	10 Days
London	UKT	07:30:00	17:00:00	10 Days
New York	EST	08:00:00	14:00:00	12 Days

### 11.2.1 The Trading Model and Its Data-Processing Environment

As in forecasting or other applications, trading models depend heavily on the quality of the financial data that are fed to the program. Problems related to bad or invalid data can play an important role at different stages of the decision process. For instance, bad data can disturb the computation of the model indicators and then imply a partial or complete loss of the prediction power related to these indicators. One other sensitive part is the computation of the current return of the open positions. The current return is often used to trigger stop-profit algorithms, or exit specific positions. Any invalid price that passes through the filter incorrectly can produce a long-term perturbation effect on the trading abilities of the system, especially if it is used as a transaction price. To avoid data-related problems, a good trading system must include a special filter to cancel or postpone recommendations until a realistic transaction price is selected.

*Trading Hours and Market Holidays* Although some markets like the FX market operate continuously, individual traders or institutions generally partake of this market only for a portion of each day. Our models accommodate such users by incorporating the notion of business hours and holidays. Every trading model is associated with a local market that is identified with a corresponding geographical region. In turn, this is associated with generally accepted office hours and public holidays. The local market is defined to be open at any time during office hours provided that the trading model does not operate on a weekend or a public holiday. Typical opening hours for a model are between 9:00 and 17:00 local time, the exact times depending on the particular local market and traded instruments. In the case of FX, Table 11.1 presents typical opening hours of different geographical markets. Except for closing an open position if the price hits a stop loss limit (described in section 11.2.1), a model may not deal outside of the market's opening hours.

TABLE 11.1 Market constraints.

**Current Return Calculations** In the trading room, people seldom take a full exposure at once. The traders like to build their positions in steps (gearing steps). In such cases, it is useful to introduce an auxiliary variable, the average price  $\bar{p}$  paid for achieving the current exposure (gearing). This variable simplifies the computation of the return of a position built in steps. After a new deal with index  $i$ , the average price depends on the type of transaction as follows:

$$\bar{p}_i \equiv \begin{cases} \bar{p}_{i-1} & \text{if } |g_i| < |g_{i-1}| \text{ and } g_i g_{i-1} > 0 \\ g_i \left[ \frac{g_i - g_{i-1}}{p_i} + \frac{g_{i-1}}{\bar{p}_{i-1}} \right]^{-1} & \text{if } |g_i| > |g_{i-1}| \text{ and } g_i g_{i-1} > 0 \\ p_i & \text{if } g_i g_{i-1} < 0 \text{ or } g_{i-1} = 0 \\ \text{undefined} & \text{if } g_i = 0 \end{cases} \quad (11.1)$$

where  $g_{i-1}$  and  $g_i$  are the previous and current gearings, respectively,  $p_i$  is the current transaction price, and  $\bar{p}_{i-1}$  is the average price before the deal. In the initial case, when the current gearing is neutral, the average price  $\bar{p}$  is not yet defined. If we start from a neutral position  $g_{i-1} = 0$  or reverse a position  $g_i g_{i-1} < 0$ , the price to build the position is simply the current price  $p_i$ . If the new position is built on top of a previous position, then we need to compute the average price paid from the price paid for each fraction of the full position. If the new position is just unfolding part of the previous position, then the average price paid for the position does not change. It is simply either profit taking or stop loss.

The average price  $\bar{p}$  is needed to compute a quantity central to a trading model, the *return* of a deal,

$$r_i \equiv (g_{i-1} - g'_i) \left( \frac{p_i}{\bar{p}_{i-1}} - 1 \right) \quad (11.2)$$

where the gearing  $g'_i$  is equal to 0 if the model takes an opposite position ( $g_i g_{i-1} < 0$ ) and  $g_i$  otherwise. There are deals with no return: those starting from a neutral gearing,  $g_{i-1} = 0$ , and those increasing the absolute value of the gearing while keeping its sign.<sup>2</sup> In these cases, Equation 11.2 does not apply (whereas Equation 11.1 applies to all deals).

The *current return*,  $r_c$ , is the unrealized return of a transaction when the current position is off the equilibrium ( $g_i \neq 0$ ). If  $p_c$  is the current market price required

<sup>2</sup> The example below demonstrates the accounting of a trading model of USD-CHF, where CHF is the *home (numeraire)* currency and USD is the *foreign exchanged* currency. The trading model is played with a limit of 100 CHF. The usual practice for the capital flow in foreign exchange trading is that it is started from a capital of zero with credit limit. This is what is assumed here. All of our return calculations are expressed in terms of the home currency. In other words, the returns are calculated in terms of DEM for USD-DEM, CHF for USD-CHF, FRF for USD-FRF, and JPY for DEM-JPY.

for going back to neutral, generalizing Equation 11.2 yields the current return,

$$r_c \equiv g_i \left( \frac{p_c}{p_i} - 1 \right) \quad (11.3)$$

**Gearing Calculation** A gearing calculator lies at the heart of a trading model. The gearing calculator provides the trading model with its intelligence and the ability to capitalize on movements in the markets. The gearing calculator also provides the trading model with particular properties. These include the frequency of dealing and the circumstances under which positions may be entered.

In other words, the gearing calculator is the *real* model. In contrast, the other trading model components form a shell around the gearing calculator, providing it with price data, detecting if the stop-loss is hit, and examining the trading recommendations made by the gearing calculator. The gearing calculator reevaluates its position every time a new price quote is received from the data-vendors. (As previously noted, a filter validates each price beforehand in order to eliminate outliers and other implausible data.)

The gearing calculator employs two kinds of ingredients: a set of indicators, which are produced from the input price data, and trading rules, which are functions of the past dealing history, the current position, and other quantities such as the current unrealized return of an open position.

The models described here give a recommendation not only for the direction but also for the amount of the exposure. In our models, the possible exposures (gearings) are  $\pm\frac{1}{2}$ ,  $\pm 1$  (full exposure) or 0 (no exposure).

**Recommendation Maker** The fact that the gearing calculator's indicators and rules suggest entering a new position does not necessarily mean that the model will make such a recommendation. Whether it does or not depends on various secondary rules that then take effect.

These rules constitute the deal acceptor. This determines whether the deal proposed by the indicators is allowed to be made. The prime constraint is the timing of the proposed deal. First, no deal other than a stop-loss deal (see Section 11.2.1) may take place within a few minutes of a deal already having occurred. This is to prevent overloading a human dealer who may be following the models. Second,

Time	Gearing	Current position in CHF	Current position in USD	FX
0	0	0	0	Don't care
1	0.5	-50	35.71	1.4
2	1	-100	69.04 (33.33 more)	1.5
3	0	10.46 (69.04*1.6 more)	0	1.6

In the example above, the trading lots in CHF are always 50 (half gearing step) or 100 (full gearing step) when increasing the (long or short) position, whereas decreasing the position means selling the full current USD amount (when going to neutral) or half the current USD amount (when going from gearing 1 to 1/2). There can be other accounting conventions, but they hardly differ numerically.

the gearing calculator may make a recommendation to enter a new trading position but this recommendation can be followed only if the local market is open.

The quality of the most recent price imposes another constraint. A stringent filter determines if a given price is suitable for dealing. This is to ensure that recommended deals are made only with genuine prices rather than extraneous data. The deal acceptor permits a new deal only with a price passing the deal-filter.

If the gearing calculator suggests entering a new position but the deal acceptor decrees otherwise, the suggestion is simply ignored. Eventually, when timing and other factors are right, the gearing calculator will suggest entering a new position and the deal acceptor will approve.

*Stop-Loss Detection* Besides being passed on to the gearing calculator, the filtered price quotes are also sent to the stop-loss detector. The stop-loss detector is triggered if the market moves in an unexpected direction. That is, if the model enters a trading position because it anticipates the market to move in a certain direction but in fact the market then moves the other way, the stop-loss may be hit. The trading model defines a stop-loss price when a position is entered. If the current price – that is, the most recent price – moves below the stop-loss price (for a long position) or above the stop-loss price (for a short position), the stop-loss is said to be hit. Hitting the stop-loss causes a deal to close the current open position (i.e., return to the neutral position). In effect, the stop-loss prevents excessive loss of capital when the market moves in an unexpected direction. The stop-loss price may change when a new position is entered or as the current price changes (see Section 11.2.1). The current stop-loss price is displayed on the user-agent.

For 24-hr markets like FX, a stop-loss deal may occur at any time, even outside local market hours. In this case, the assumption is that a position that is kept open outside market hours is handled by a colleague present in another market-place who will deal appropriately if the stop-loss is hit. Should this happen, no further change in position occurs until the local market opens once again.

*Stop-Profit Control* The concept of stop-profit is associated with that of stop-loss. The stop-loss price starts to move in parallel with the current price once a trading model has achieved a potential profit (3% or slightly less in FX market) since entering the latest position. In other words, being in a situation whereby the model could realize such a gain by immediately entering a neutral position causes the stop-loss price to start moving. The difference between the stop-loss and current prices is kept constant as long as the current price continues moving in a direction that increases the potential profit of the open position. That is, the stop-loss price moves as a ratchet in parallel with the current price. The stop price is allowed to move only during opening hours. It is never adjusted when the market is closed.

The model then enters a neutral position if it detects prices slipping backward. This allows a model to save any profit it has generated rather than lose it when the market abruptly turns. This one-directional movement of the stop-loss price allows the model to capitalize on a price trend.

T O R O N T O



### 11.2.2 Simulated Trader

The simulated trader allows the system to control continuously its performance by simulating a trade every time the trading model gives a recommendation. In the following, we shall describe the different part that composes the simulated trader.

*Opportunity Catcher* The trading model may make a deal recommendation in two distinct ways. One, the gearing calculator may make a recommendation that is then authorized by the deal acceptor. Two, hitting the stop-loss price activates the stop-loss detector.

Whichever way a deal comes about, the *opportunity catcher* is activated. The opportunity catcher manifests itself on the user-agent as an eye-catching signal for the FX dealer to buy or sell according to the recommendation.

While he or she is actively dealing, the opportunity catcher in the trading model collects the transaction price with which to deal, either the median bid price if going from a longer position to a shorter one or the median ask price if going from a shorter position to a longer one. This search for the transaction price lasts for 2 or 3 min depending on the currency, the assumption being that a quoted price has a life-time of about 2 or 3 min even if it is superseded by later quotes.

After the 2 or 3 min search period, a second signal appears on the user-agent signifying that the trading model has made a simulated deal using the transaction price found by the opportunity catcher. The FX dealer then concludes his or her deal-making activities and waits until the trading model produces another recommendation.<sup>3</sup>

*Bookkeeper* The *bookkeeper* executes simulated deals on behalf of the trading model. It keeps track of all deals that have been made and evaluates statistics demonstrating the performance of the trading model. The bookkeeper computes a set of quantities that are important for the different trading rules like the following:

- The *maximum return when open*, which is the maximum value of  $r_c$  from a transaction  $i$  to a transaction  $i + 1$  reached during opening hours,
- The *minimum return when open*, which is the minimum value of  $r_c$  from a transaction  $i$  to a transaction  $i + 1$  reached during opening hours.

In this section we describe some of the important variables that need to be watched for deciding on the quality of a specific model. These are the following:

- The *total return*,  $R_T$ , is a measure of the overall success of a trading strategy over a period  $T$ , and is defined by

$$R_T \equiv \sum_{j=1}^n r_j \quad (11.4)$$

<sup>3</sup> As a point of detail, the opportunity catcher is not activated for a stop-loss deal occurring outside market hours. In this case the trading model deals directly. A human trader following the model should then make a corresponding deal for himself as quickly as possible.

where  $n$  is the total number of transactions during the period  $T$ ,  $j$  is the  $j^{\text{th}}$  transaction and  $r_j$  is the return from the  $j^{\text{th}}$  transaction. The total return expresses the amount of profit (or loss) made by a trader always investing up to his/her initial capital or credit limit in his/her home currency.

- The *cumulated return*,  $C_T$ , is another measure of the overall success of a trading model wherein the trader always reinvests up to his/her current capital, including gains or losses

$$C_T \equiv \prod_{j=1}^n (1 + r_j) - 1 . \quad (11.5)$$

This quantity is slightly more erratic than the total return.

- The *maximum drawdown*,  $D_T$ , over a certain period  $T = t_E - t_0$ , is defined by

$$D_T \equiv \max( R_{t_a} - R_{t_b} \mid t_0 \leq t_a \leq t_b \leq t_E ) \quad (11.6)$$

where  $R_{t_a}$  and  $R_{t_b}$  are the total returns of the periods from  $t_0$  to  $t_a$  and  $t_b$ , respectively.

- The *profit over loss ratio* provides information on the type of strategy used by the model. Its definition is

$$\frac{P_T}{L_T} \equiv \frac{N_T(r_j \mid r_j > 0)}{N_T(r_j \mid r_j < 0)} \quad (11.7)$$

where  $N_T$  is a function that gives the number of elements of a particular set of variables under certain conditions during a period  $T$ . Here the numerator corresponds to the number of profitable deals over the period  $T$  and the denominator is the number of losing deals over the same period.

### 11.3 RISK SENSITIVE PERFORMANCE MEASURES

Evaluating the performance of an investment strategy generally gives rise to many debates. This is due to the fact that the performance of any financial asset cannot be measured only by the increase of capital but also by the risk incurred during the time required to reach this increase. Returns and risk must be evaluated together to assess the quality of an investment. In this section we describe the various performance measures used to evaluate trading models.

The annualized return,  $\bar{R}_{T,A}$ , is calculated by multiplying the total return (Equation 11.4) with the ratio of the number of days in a year to the total number of days in the entire period.<sup>4</sup> In order to achieve a high performance and good

<sup>4</sup> If it is the annualization of one particular return (for one trade going from neutral to neutral), one simply needs to multiply the return by the ratio of 1 year in days to the time interval from neutral to neutral. Usually, the annualization of the total return is calculated for all the trades during a whole year. This is simply the sum of all trade returns not annualized during the whole year. If at the end of the year there is an open position, the current return of your open position is added to the total return.

acceptance among investors, investment strategies or trading model performance should provide high annualized total return, a smooth increase of the equity curve over time, and a small clustering of losses. The fulfilment of these conditions would account for a high return and low risk. In addition to favoring this type of behavior, a performance measure should present no bias toward low-frequency models by including always the unrealized return of the open position and not only the net result after closing the position.

Already in 1966, Sharpe (1966) introduced a measure of mutual funds performance, which he called at that time a *reward-to-variability ratio*. This performance measure was to later become the industry standard in the portfolio management community under the name of the Sharpe ratio, Sharpe (1994). Practitioners frequently use the Sharpe Ratio to evaluate portfolio models. The definition of the Sharpe ratio is

$$S_T \equiv A_{\Delta t} \frac{\bar{r}}{\sqrt{\sigma_r^2}} \quad (11.8)$$

where  $\bar{r}$  is the average return and  $\sigma_r^2$  is the variance of the return around its mean and  $A_{\Delta t}$  is an annualization factor,<sup>5</sup> depending on the frequency at which the returns are measured Sharpe (1994).<sup>6</sup> Unfortunately, the Sharpe ratio is numerically unstable for small variances of returns and cannot consider the clustering of profit and loss trades.

There are many aspects to the trading model performance; therefore, different quantities have to be computed to assess the quality of a model. In the section on the bookkeeper, we already described some of the important variables that need to be watched for deciding on the quality of a specific model. Here we introduce the two risk-sensitive measures that are the basic quantities used in further sections to analyze the behavior of trading models.

### 11.3.1 $X_{eff}$ : A Symmetric Effective Returns Measure

As the basis of a risk-sensitive performance measure, we define a cumulative variable  $\tilde{R}_t$ , at time  $t$ , as the sum of the total return  $R_T$  of Equation 11.4 and the unrealized current return  $r_c$  (Equation 11.3) of the open position. This quantity reflects the current value of the investment and includes not only the results of previously closed transactions but also the value of the open position (mark-to-market). This means that  $\tilde{R}_t$  is measuring the risk independently of the actual trading frequency of the model. Similar to the difference between price and returns, the variable of relevance for the utility function is the change of  $\tilde{R}$  over a time

<sup>5</sup>  $A_{\Delta t} = \sqrt{12}$  for monthly frequency.

<sup>6</sup> Here the Sharpe ratio refers to the calculation of the returns in the expressed currency and the variance is computed with monthly returns. The monthly returns are the total return achieved to the end of the month (sum of all returns up to now, including the current return of the open position) minus the total return achieved at the end of the previous month.

interval  $\Delta t$ ,

$$X_{\Delta t} = \tilde{R}_t - \tilde{R}_{t-\Delta t} \quad (11.9)$$

where  $t$  expresses the time of the measurement. Generally  $\Delta t$  is allowed to vary from 7 days to 301 days. A risk-sensitive measure of trading model performance can be derived from the utility function framework (Keeney and Raiffa, 1976). Let us assume that the variable  $X_{\Delta t}$  follows a Gaussian random walk with mean  $\bar{X}_{\Delta t}$  and the risk aversion parameter  $\alpha$  is constant with respect to  $X_{\Delta t}$ . The resulting utility  $u(X_{\Delta t})$  of an observation is  $-\exp(-\alpha X_{\Delta t})$ , with an expectation value of  $\bar{u} = u(\bar{X}_{\Delta t}) \exp(\alpha^2 \sigma_{\Delta t}^2 / 2)$ , where  $\sigma_{\Delta t}^2$  is the variance of  $X_{\Delta t}$ . The expected utility can be transformed back to the *effective return*,  $X_{eff} = -\log(-\bar{u})/\alpha$  where

$$X_{eff} = \bar{X}_{\Delta t} - \frac{\alpha \sigma_{\Delta t}^2}{2} \quad (11.10)$$

The risk term  $\alpha \sigma_{\Delta t}^2 / 2$  can be regarded as a risk premium deducted from the original return where  $\sigma_{\Delta t}^2$  is computed by

$$\sigma_{\Delta t}^2 = \frac{n}{n-1} \left( \overline{X_{\Delta t}^2} - \bar{X}_{\Delta t}^2 \right) \quad (11.11)$$

Unlike the Sharpe ratio, this measure is numerically stable and can differentiate between two trading models with a straight-line behavior ( $\sigma_{\Delta t}^2 = 0$ ) by choosing the one with the better average return.<sup>7</sup> The measure  $X_{eff}$  still depends on the size of the time interval  $\Delta t$ . It is hard to compare  $X_{eff}$  values for different intervals. The usual way to enable such a comparison is through the annualization factor,  $A_{\Delta t}$ , where  $A_{\Delta t}$  is the ratio of the number of  $\Delta t$  in a year divided by the number of  $\Delta t$ 's in the full sample

$$X_{eff,ann,\Delta t} = A_{\Delta t} X_{eff} = \bar{X} - \frac{\alpha}{2} A_{\Delta t} \sigma_{\Delta t}^2 \quad (11.12)$$

where  $\bar{X}$  is the annualized return and it is no longer dependent on  $\Delta t$ . The factor  $A_{\Delta t} \sigma_{\Delta t}^2$  has a constant expectation, independent of  $\Delta t$ . This annualized measure still has a risk term associated with  $\Delta t$  and is insensitive to changes occurring with much longer or much shorter horizons. To achieve a measure that simultaneously considers a wide range of horizons, a weighted average of several  $X_{eff,ann}$  is computed with  $n$  different time horizons  $\Delta t_i$ , and thus takes advantage of the fact that annualized  $X_{eff,ann}$  can be directly compared,

$$X_{eff} = \frac{\sum_{i=1}^n w_i X_{eff,ann,\Delta t_i}}{\sum_{i=1}^n w_i} \quad (11.13)$$

<sup>7</sup> An example for the limitation of the Sharpe ratio is its inability to distinguish between two straight line equity curves with different slopes.

where the weights  $w$  are chosen according to the relative importance of the time horizons  $\Delta t_i$  and may differ for trading models with different trading frequencies. Generally,  $\alpha$  is set to  $\alpha = 0.1$  when the returns are expressed as a percentage. If they are expressed in numbers,  $\alpha$  would be equal 10. The risk term of  $X_{eff}$  is based on the volatility of the total return curve against time, where a steady, linear growth of the total return represents the zero volatility case. This volatility measure of the total return curve treats positive and negative deviations symmetrically, whereas foreign exchange dealers become more risk averse in the loss zone and hardly care about the clustering of positive profits.

### 11.3.2 $R_{eff}$ : An Asymmetric Effective Returns Measure

A measure that treats the negative and positive zones asymmetrically is defined to be  $R_{eff}$ , (Müller *et al.*, 1993b; Dacorogna *et al.*, 2001b) where  $R_{eff}$  has a high risk aversion in the zone of negative returns and a low one in the zone of profits, whereas  $X_{eff}$  assumes constant risk aversion. A high risk aversion in the zone of negative returns means that the performance measure is dominated by the large drawdowns. The  $R_{eff}$  has two risk aversion levels: a low one,  $\alpha_+$ , for positive  $\Delta \tilde{R}_t$  (profit intervals) and a high one,  $\alpha_-$ , for negative  $\Delta \tilde{R}_t$  (drawdowns),

$$\alpha = \begin{cases} \alpha_+ & \text{for } \Delta \tilde{R}_t \geq 0 \\ \alpha_- & \text{for } \Delta \tilde{R}_t < 0 \end{cases} \quad (11.14)$$

where  $\alpha_+ < \alpha_-$ . The high value of  $\alpha_-$  reflects the high risk aversion of typical market participants in the loss zone. Trading models may have some losses but, if the loss observations strongly vary in size, the risk of very large losses becomes unacceptably high. On the side of the positive profit observations, a certain regularity of profits is also better than a strong variation in size. However, this distribution of positive returns is never as vital for the future of market participants as the distribution of losses (drawdowns). Therefore,  $\alpha_+$  is smaller than  $\alpha_-$  and we assume that  $\alpha_+ = \alpha_-/4$  and  $\alpha_- = 0.20$ . These values are under the assumption of the return measured as *percentage*. They have to be multiplied by 100 if the returns are not expressed as percentage figures.

The *risk aversion*  $\alpha$  associated with the utility function  $u(\Delta \tilde{R})$  is defined in Keeney and Raiffa (1976) as follows:

$$\alpha = - \frac{\frac{d^2 u}{[d(\Delta \tilde{R})]^2}}{\frac{du}{d(\Delta \tilde{R})}} \quad (11.15)$$

The utility function is obtained by inserting Equation 11.14 in Equation 11.15 and integrating twice over  $\Delta \tilde{R}$ :

$$u = u(\Delta \tilde{R}) = \begin{cases} -\frac{e^{-\alpha_+ \Delta \tilde{R}}}{\alpha_+} & \text{for } \Delta \tilde{R} \geq 0 \\ \frac{1}{\alpha_-} - \frac{1}{\alpha_+} - \frac{e^{-\alpha_- \Delta \tilde{R}}}{\alpha_-} & \text{for } \Delta \tilde{R} < 0 \end{cases} \quad (11.16)$$

The utility function  $u(\Delta \tilde{R})$  is monotonically increasing and reaches its maximum 0 in the case  $\Delta \tilde{R} \rightarrow \infty$  (infinite profit). All other utility values are negative. (The absolute level of  $u$  is not relevant; we could add and/or multiply all  $u$  values with the same constant factor(s) without affecting the essence of the method.)

The inverse formula computes a return value from its utility:

$$\Delta \tilde{R} = \Delta \tilde{R}(u) = \begin{cases} -\frac{\log(-\alpha_+ u)}{\alpha_+} & \text{for } u \geq -\frac{1}{\alpha_+} \\ -\frac{\log(1 - \frac{\alpha_-}{\alpha_+} - \alpha_- u)}{\alpha_-} & \text{for } u < -\frac{1}{\alpha_+} \end{cases} \quad (11.17)$$

The more complicated nature of the new utility definition, Equation 11.16, makes deriving a formula for the *mean* utility quite difficult and offers no analytical solution. Moreover, the  $R_{eff}$  is dominated by the drawdowns that are in the *tail* of the distribution, not in the center. The assumption of a Gaussian distribution, which may be acceptable for the distribution as a whole, is insufficient in the tails of the distribution, where the stop-loss, the leptokurtic nature of price changes, and the clustering of market conditions such as volatility cause very particular forms of the distribution.

Therefore, the use of *explicit* utilities is suggested in the  $R_{eff}$  algorithm. The end results,  $R_{eff}$  and the effective returns for the individual horizons, will however, be transformed back with the help of Equation 11.17 to a return figure directly comparable with the annualized return and  $X_{eff}$ . The utility of the  $j^{th}$  observation for a given time interval  $\Delta t$  is

$$u_{\Delta t, j} = u(\tilde{R}_t, -\tilde{R}_{t, -\Delta t}) \quad (11.18)$$

The total utility is the sum of the utility for each observation

$$u_{\Delta t} = \frac{\sum_{j=1}^{N_j} v_j u_{\Delta t, j}}{\sum_{j=1}^{N_j} v_j} \quad (11.19)$$

In this formula,  $N_j$  is the number of observed intervals of size  $\Delta t$  that overlaps with the total sampling period of size  $T$  and the weight  $v_j$  is the ratio of the amount of time during which the  $j^{th}$  interval coincides with the sampling period over its interval size  $\Delta t$ . This weight is generally equal to one except for the first observation(s), which can start before the sample starts, and the last one(s), which can end after the sample ends. To obtain a lower error in the evaluation of the mean utility, different regular series of overlapping intervals of size  $\Delta t$  can be used. The use of overlapping intervals is especially important when the interval size  $\Delta t$  is large compared to the full sample size  $T$ . Another argument for overlapping is the high, overproportional impact of *drawdowns* on  $R_{eff}$ . The higher the overlap factor, the higher the precision in the coverage of the worst drawdowns.

FOR OFFICIAL USE ONLY

The mean utility  $u_{\Delta t}$  can be transformed back to an *effective* return value by applying Equation 11.17:

$$\Delta \tilde{R}_{eff, \Delta t} = \Delta \tilde{R}(u_{\Delta t}) \quad (11.20)$$

This  $\Delta \tilde{R}_{eff, \Delta t}$  is the typical, effective return for the horizon  $\Delta t$ , but it is not yet annualized. As in the case of the  $X_{eff}$ , an annualization is necessary for a comparison between  $R_{eff}$  values for different intervals. The annualization factor,  $A_{\Delta t}$ , is the ratio of the number of  $\Delta t$  in a year divided by the number of nonoverlapping  $\Delta t$ 's in the full sample of size  $T$ . We have

$$\tilde{R}_{eff, ann, \Delta t} = A_{\Delta t} \Delta \tilde{R}_{eff, \Delta t} \quad (11.21)$$

To achieve a measure that simultaneously considers a wide range of horizons, we define  $R_{eff}$  as a weighted mean over all the  $n$  horizons,

$$R_{eff} = \frac{\sum_{i=1}^n w_i R_{eff, ann, \Delta t_i}}{\sum_{i=1}^n w_i} \quad (11.22)$$

where the weights  $w_i$  are chosen according to the relative importance of the time horizons  $\Delta t_i$  and may differ for trading models with different trading frequencies. In the case of the trading models described in this book, we have chosen a weighting function

$$w_i = w(\Delta t_i) = \frac{1}{2 + \left(\log \frac{\Delta t_i}{90 \text{ days}}\right)^2} \quad (11.23)$$

with the maximum for a  $\Delta t$  of 90 days.

Both  $X_{eff}$  and  $R_{eff}$  are quite natural measures. They treat risk as a discount factor to the value of the investment. In other words, the performance of the model is discounted by the amount of risk that was taken to achieve it. In the  $X_{eff}$  case the risk is treated similarly both for positive or negative outcome, whereas in the case of  $R_{eff}$  negative performance is more penalized.

#### 11.4 TRADING MODEL ALGORITHMS

We turn now to the description of the techniques used to build the real-time trading models. Trading models have been developed for many decades by a large number of people and applied in all types of financial markets. These models have been designed from a broad class of indicators ranging from classical technical analysis up to chaotic theory. It is just not possible to provide a comprehensive list of the various approaches used in trading system design and the literature in this field is so large that we will leave this exercise to other authors. As in the case of technical analysis, hundreds of articles and books have been written.

Generally, trading systems are built from a few classes of indicators providing specific types of information on the underlying financial time series. For instance, we have these:

- The *trend following* indicators, which allow to detect and follow major market trends.
- The *overbought and oversold* indicators, which allow to detect important market turning points.
- The *cycle* indicators, which try to emphasize periodic market fluctuations.
- The *timing* indicators, which provide optimum exit conditions.

As an example we will describe one model that we have developed for the FX market and that many large banks have actively used for a decade. As we have pointed out earlier, a trading strategy is built from some indicators and a set of decision rules. Indicators are variables of the trading system algorithm whose values, together with the system rules, determine the trading decision process. In Chapter 3.3 we gave different descriptions of indicators that have been used in conjunction with trading models.

#### 11.4.1 An Example of a Trading Model

The real-time trading model (RTT) studied in this section is classified as a one-horizon, high-risk/high-return model. The RTT is a trend-following model and takes positions when an indicator crosses a threshold. The indicator is momentum based, calculated through specially weighted moving averages with repeated application of the exponential moving average operator (see Section 3.3.6). In the case of extreme foreign exchange movements, however, the model adopts an overbought/oversold (contrarian) behavior and recommends taking a position against the current trend. The contrarian strategy is governed by rules that take the recent trading history of the model into account. The RTT model goes neutral only to save profits or when a stop-loss is reached. Its profit objective is typically at 3%. When this objective is reached, a gliding stop-loss prevents the model from losing a large part of the profit already made by triggering it to go neutral when the market reverses.

At any point in time  $t$ , the gearing function for the RTT is

$$g_t(I_x) = \text{sign}(I_x(t)) f(|I_x(t)|) c(I(t))$$

where

$$I_x(t) = x_t - MA(\tau = 20 \text{ days}, 4; x)$$

where  $x_t$  is the logarithmic price at time  $t$  and the moving average (MA) of  $x$  follows the definition and notation of Equation 3.56 (where the last argument  $x$  of



MA indicates the time series to which the MA operator is applied),

$$f(|I_x(t)|) = \begin{cases} \text{if } |I_x(t)| > b & 1 \\ \text{if } a < |I_x(t)| < b & 0.5 \\ \text{if } |I_x(t)| < a & 0 \end{cases}$$

and

$$c_t(I) = \begin{cases} +1 & \text{if } |I_x(t)| < d \\ -1 & \text{if } |I_x(t)| > d \text{ and } g_{t-1} \cdot \text{sign}(I_x(t)) > 0 \text{ and } r_l > P \end{cases}$$

where  $a < b < d$  and  $r_l$  is the return of the last deal and  $P$  the profit objective. The function,  $f(|I_x(t)|)$ , measures the size of the signal at time  $t$  and the function,  $c(|I_x|)$ , acts as a contrarian strategy. The model will enter a contrarian position only if it has reached its profit objective with a trend following position. In a typical year, the model will play against the trend two to three times while it deals roughly 60 to 70 times. The hit rate of the contrarian strategy is of about 75%.

The parameters  $a$  and  $b$  depend on the position of the model,

$$a(t) = \begin{cases} a & \text{if } g_{t-1} \neq 0 \\ 2a & \text{if } g_{t-1} = 0 \end{cases}$$

and  $b = 2a$ . The thresholds are also changed if the model is in a position  $g_t \neq 0$  and the volatility of the price has been low, in the following way:

$$a(t) = \begin{cases} a & \text{if } |x_e - x_t| > v \\ 10a & \text{if } |x_e - x_t| < v \end{cases}$$

where  $x_e$  is the logarithmic entry price of the last transaction and  $v$  is a threshold, generally quite low  $< 0.5\%$ . This means that the model is only allowed to change position if the price has significantly moved from the entry point of the deal.

Because  $X_{eff}$  and  $R_{eff}$  are implicit functions of the gearing, the optimization of the RTT model is based on the  $X_{eff}$  and  $R_{eff}$  performance. The parameters subject to optimization are;  $\tau$ ,  $a$ ,  $d$ , and  $v$ . There are two other auxiliary parameters, which are the stop loss,  $S$ , at which an open position is automatically closed and the profit objective,  $P$ . These parameters are only optimized at the end once the others have been found and they are also not allowed to vary all the way because maximum stop-loss and maximum gain limits are set by the environment.<sup>8</sup> The model is subject to the open-close and holiday closing hours of the Zurich market.

#### 11.4.2 Model Design with Genetic Programming

The major problem with trading models is the large amount of time needed to develop and optimize new trading strategies. As we said before, a trading strategy is a small computer program composed of some indicators to forecast price trends combined with a set of rules to determine the trading decision process.

<sup>8</sup> For more details on the optimization procedure, see Pictet *et al.* (1992).

One very promising approach in the search of new trading strategies is provided by genetic programming (GP) method (Koza, 1992; Banzhaf *et al.*, 1998). This is an evolutionary algorithm that allows to automatically discover computer programs that solve a given problem. Evolutionary algorithms tend to find globally satisfactory solutions to the problem and, much in the same way as in nature, populations of organisms tend to adapt to their surrounding environment. Such an approach has been applied to stock indices, as noted in Allen and Karjalainen (1999), and to exchange rates, as noted in Oussaidène *et al.* (1997); Neely *et al.* (1997); Bhattacharyya *et al.* (1998).

Individual programs in GP are represented as parse trees with ordered branches in which the internal nodes are functions (with subtree branches as function arguments) and the leaves or terminal nodes are variables. The functions are chosen from a user-defined *function set*, those that are a priori believed to be useful for the problem at hand, and the leaves are selected from a *terminal set* containing the principal variables or constants of the problem.

Once an initial population has been created, the genetic algorithm enters a loop. At the end of each iteration (or generation), a new population has been created by applying a certain number of stochastic operators to the members of the previous population. A *selection* operator is first applied in order to extract some above-average individuals for reproduction. When a population of parents has been extracted, two reproduction operators are used: *crossover* and *mutation*. As shown on Figure 11.2, the crossover operator starts by selecting a random crossover point in each parent tree (*a* and *b*) and then exchanging the subtrees, giving rise to two offspring trees (*c* and *d*). The crossover sites, *c1* and *c2*, are usually chosen with nonuniform probability, in order to favor internal nodes with respect to leaves. In the same figure we can observe that from two parents trees, which are not interesting, the crossover is able to generate the offspring (*c*), which correspond to a well-known simple trading strategy using the difference between two moving averages. After crossover, a certain proportion of the offspring are subject to mutation. The mutation operator is implemented by randomly removing a subtree at a selected point and replacing it with a randomly generated subtree.

In the basic genetic programming approach, it is generally required that all elements of a tree return the same data type, so as to allow arbitrary subtrees to be recombined by the crossover and mutation operators. This *closure property* (Koza, 1992) can be a potential limitation in some applications like the trading strategies search process. In earlier studies, on the use of GP for searching new trading strategies (Oussaidène *et al.*, 1997), a large proportion of the population members were noted to be irrelevant and resulted in a wasteful search. For instance, if you carefully study the different GP trees in Figure 11.2, you can easily conclude that only the offspring (*c*) corresponds to a reasonable indicator.

As mentioned, trading models are a function of the price history. In the case of the FX market, it is common to consider the logarithmic middle price  $x$ , which possesses an exact symmetry  $x \rightarrow -x$ , corresponding to the interchange of the expressed and exchanged currencies. Consider the U.S. Dollar to German Mark

E  
N  
G  
L  
I  
S  
H  
L  
A  
N  
G  
U  
A  
G  
E




FIGURE 11.2 Crossover operator.

(USD-DEM) exchange rate. A trading model that is optimum for the USD-DEM rate is also expected to be optimum for the inverted DEM-USD rate. For this to hold, the model for the inverted DEM-USD rate should provide a signal, at any time  $t$ , that has exactly the reverse sign than the one for the USD-DEM rate. It means the output signal  $g_t(x)$  of a consistent trading model must be an antisymmetric function of the return. Enforcing the symmetry condition on the trading model thus requires that  $g_t(x) \rightarrow g_t(-x) = -g_t(x)$ . Maintaining this property in a GP tree requires tracking the symmetry property at individual nodes in the tree, and forms the basis for defining syntactic restrictions. To enforce the symmetry, three possible types are defined:

- Antisymmetric type A (e.g., a moving average of  $x$ ):  
 $A(-x) = -A(x)$
- Symmetric type S (e.g., a volatility):  $S(-x) = S(x)$
- Constant type C (numerical constants)

Constants are in essence of symmetric type; however, there are many advantages to considering the constant type separately. In specifying a GP trading model,

Downloaded from https://www.cambridge.org/core. IP: 128.112.1.104, on 05 Nov 2020 at 10:00:00, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9780511526351.012

+ -	A	S	C
A	A	-	-
S	-	S	S
C	-	S	-

* /	A	S	C
A	S	A	A
S	A	S	S
C	A	S	-

FIGURE 11.3 Syntactic restrictions for basic arithmetic operators.

every node evaluation is considered to return both a value and a type (A, S, or C as defined above). The typing mechanism is used to categorize the symmetry properties. A variety of functions may be considered for formulating trading models. Each function must be specified in terms of syntactic restrictions relating to symmetry, and guiding their combination with terminals and other functions. As an example, the syntactic restrictions for the basic arithmetic operators are provided in Figure 11.3. In these tables, the first row and column correspond to the type of the two arguments, with the intersection cells showing the type of the result. The symbol “-” represents a combination of arguments that is disabled. Operation on two constants is generally disabled to avoid wasteful computation of constants through the regular crossover and mutation operators (Evet and Fernandez, 1998). The constants are generally mutated using specific non-uniform mutation operators (Michalewicz, 1994).

Other classes of functions are useful in the construction of the trading strategies. The important one is the class of the time series convolution operators described in Section 3.3. For instance, in Figure 11.2, we see exponential moving averages of the middle logarithmic price  $x$  for the two ranges of 20 and 40 days. Such operators can be used as function nodes if their syntactic restrictions are provided. They can also be used directly as input variables of the antisymmetric or symmetric type, depending on their symmetry properties when  $x$  is replaced by  $-x$ . In the case of a moving average operator, syntactic restrictions are forced to have the same symmetry type for the input and output signals.

Other operators often used in trading strategies are the classes of:

- Logic operators: AND, OR,
- Comparison operators: greater than, smaller than,
- Conditional operators: IF-then, IF-then-else.

As we expect an asymmetric output trading signal, it is suitable to use for these operators a ternary Boolean logic, Oussaidène *et al.* (1997).

In using such trading strategies, like in Bhattacharyya *et al.* (1998), we clearly depart from the closure property. In fact, we consider here a strongly typed GP approach (see Montana, 1995), where the evolution procedure needs to define a random tree initialization routine, and crossover and mutation operators respecting the defined restrictions.

As an example, we describe a study by Chopard *et al.* (2000) who analyze five exchange rates (USD-DEM, USD-JPY, USD-CHF, GBP-USD, and USD-FRF),

where each time series contains 9 years of hourly data from January, 1, 1987, to December, 31, 1995. The data are divided into alternate training (in-sample) and test (out-of-sample) periods of  $1\frac{1}{2}$  year each. We have used the four arithmetic functions (+, -, \*, /), the basic comparison (<, >), and the logical operators (AND, OR, IF), which were defined in Bhattacharyya *et al.* (1998). Both comparison and logical operators are defined for a ternary logic {-1, 0, +1}, which corresponds to the signal returned by a trading model. As in the case of the arithmetic functions, these operators are used with syntactic restrictions to preserve the overall symmetry properties of the generated GP trees. Indicators reported in earlier studies (Oussaidène *et al.*, 1997; Bhattacharyya *et al.*, 1998) are also used for this application. The terminals used are as follows:

- Antisymmetric indicators:  $M_n$  that represent the momentum of price  $x$  over  $n$  days. Here we consider three different ranges:  $M_8$  and  $M_{16}$ ,  $M_{32}$ .
- Symmetric indicators:  $V_n$  that are volatility indicators over  $n$  days:  $V_8$  and  $V_{16}$ .
- Constant values in the range [-2, +2].

As a last step, the output value of the GP tree is mapped to a gearing value in the range [-1, +1] to obtain a gearing signal. For the purpose of reducing overfitting, each trading model is tested on many exchange rate time series. The fitness measure,  $X_{eff}$ , of a GP-tree is then defined as the average fitness over each exchange rate, decreased by a penalty proportional to the standard deviation of these values.

Five independent optimization runs are performed. In each run, we evolve four subpopulations of 100 individuals each. All subpopulations are created randomly using the *ramped half-half* approach (Koza, 1992) with a maximum depth of 4. In the reproduction phase, we use the *tournament* selection. The mutation and crossover operators are used with corresponding probabilities of 20% and 80% and the maximum depth allowed for generated trees is fixed to 6. Each subpopulation sends periodically 5% of its best individuals to another, randomly selected, subpopulation. A given subpopulation includes its local buffer of received migrants when the difference between the best and the average scores of the current population is smaller than half the standard deviation of the scores. The new individuals replace an equal number of low-fitness ones in the receiving population. The evolution is stopped when a maximum number of 10,000 individuals have been evaluated. The selected solution is the best solution found by the four subpopulations.

Table 11.2 presents the performance results of the different runs. The first entry gives the average quality of the basic antisymmetric momentum indicators  $M_8$ ,  $M_{16}$ , and  $M_{32}$ . The results of the optimization runs are given in decreasing order of their out-of-sample performance. The results of this table indicate that on average the performance of the solutions provided by the genetic algorithm are significantly higher than the performance of the basic momentum indicators.

**TABLE 11.2** Trading model results versus tree complexity.

Tree complexity, yearly return  $R$ , and fitness value  $X_{eff}$  (in percent) corresponding to the in-sample and out-of-sample periods. The results are given for preoptimized indicators and for the best solution of each optimization run.

Run	Tree complexity	In-sample		Out-of-sample	
		$\bar{R}$	$X_{eff}$	$\bar{R}$	$X_{eff}$
Indicators	1	4.23%	-1.84%	3.61%	-3.77%
Run 1	21	8.12%	3.86%	6.29%	1.09%
Run 2	17	9.49%	4.16%	5.80%	0.47%
Run 3	19	7.45%	2.82%	5.06%	-2.34%
Run 4	24	7.71%	3.55%	4.87%	-3.28%
Run 5	28	6.84%	2.63%	3.10%	-5.85%

For instance, the solution selected in the second run that provides the best in-sample result is given by the GP-tree:

```
(IF (> V8 0.027)
  (IF (> V16 0.108)
    (+ M16 M32)
    (* M32 1.957) )
  (* M8 1.542)
)
```

The use of syntactic restrictions allows the discovery trees of lower complexity on average, compared with the previous study of Oussaidène *et al.* (1997), and all the generated GP trees are valid. However, there may exist some solutions that do not generate any trades, because some conditions always evaluate the same value. These solutions are quickly eliminated in the selection process. On average the solutions seem to be more robust and to provide higher out-of-sample performance.

One limiting problem in these optimizations was implied by the use of “hard” logical and basic comparison operators that give rise to undesirable discontinuities due to the jumps of the Boolean variables that can occur for tiny changes of the basic indicators. One better way of implementing these operators would be through the use of fuzzy-logic. A smoother transition between different logical states can probably provide better performances.

All these optimization runs are done using hourly data, for obvious efficiency reasons. Then, when such optimization is completed, the final selected solution(s) must be tested in the complete real-time trading model environment with tick-by-tick data to check its behavior and to do some fine-tuning needed for the understanding of the final users.

## 11.5 OPTIMIZATION AND TESTING PROCEDURES

When the main trading strategy has been selected, one of the most difficult tasks is to optimize the parameters present in the model and to test the different solutions in order to select the most robust trading model to be used in the real time data. The goal is to select robust solutions that have desirable generalization properties to provide satisfactory performance in the future.

In the optimization process we expect the selection of a trading model that realizes large profits from the price moves present in the time series. We assume that these profits will be maximized when the trading model catches the dynamics of the price-generation process. Unfortunately, during the optimization phase the trading model repeatedly sees the same data set and discovers how to profit from some specific price moves that could be due to some random fluctuation of the prices. This will lead a trading model to provide poor results on real-time data. Such a model is called an “overfitted” model. To minimize overfitting during optimization, a few important elements are to be present:

- A good measure of the trading model performance
- Indicator evaluation for different time series
- Large data samples
- A robust optimization technique
- Strict testing procedures

An optimization algorithm will always try to find the best solution in the parameter space. In the case of trading models, optimization of such properties is not suitable, because a solution corresponding to the best possible parameters generally corresponds to an overfitted solution. As we argued earlier, such solutions will often generate poor generalizations in a real-time trading setting. In the next section, we shall concentrate on one robust optimization technique based on the genetic algorithm approach. This method allows the selection of a group of solutions that correspond to broad regions of the parameter space where the trading performance is higher *on average*, rather than the highest.

### 11.5.1 Robust Optimization with Genetic Algorithms

The new element we want to present in this section is a way to automatize the search for improved trading models. Genetic algorithms offer a promising approach for addressing such problems (Allen and Karjalainen, 1999). Genetic algorithms consider a population of possible solutions to a given problem and evolve it according to mechanisms borrowed from natural genetic evolution: reproduction and selection. The criterion for selecting an individual is based on its fitness to the environment, or more precisely to the quality of the solution it bears. A possible solution is coded as a chromosome (gene), which is formally the data structure containing the values of the quantities characterizing the solutions.

In the framework of the present optimization, a gene will contain the indicator parameters, a time horizon, a weighting function of the past, and the type of

T  
E  
S  
T  
E  
D  
I  
N  
G  
P  
R  
O  
C  
E  
D  
U  
R  
E  
S



operations used to combine them. Contrary to GP, the gene does not offer the flexibility to the algorithm but only to the parameters. The fitness function will be based on the return obtained from the recommendations of a given trading model.

*Sharing Scheme For Multi-Modal Functions* A major problem in trading model optimization is to obtain models that are robust against market changes and random noise in data collection. In such optimization problems, sharp peaks of high fitness are usually not representative of a general solution, but rather they indicate accidental fluctuations. Such fluctuations may arise out of inherent noise in the time series or due to threshold effects in the trading model performance. Peaks in such a discontinuous, noisy and multimodal fitness space generally correspond to trading models that will not perform well in out-of-sample tests.

In the context of genetic algorithms, optimizing multimodal functions has been investigated using methods inspired from the natural notions of niche and species, as noted by Goldberg and Richardson (1987); Deb and Goldberg (1989), and Yin and Gernay (1993). The general goal is to be able to create and maintain several subpopulations, ideally one per major peak of the fitness function, instead of having the whole population converging to one global optimum.

One of the best methods was proposed by Goldberg and Richardson (1987). The idea is that the GA perception of the fitness function is changed in such a way that when individuals tend to concentrate around a high peak, the fitness is reduced by a factor proportional to the number of individuals in the region. This has the effect of diminishing the attractiveness of the peak and allowing parts of the population to concentrate on other regions. This effective fitness of an individual  $i$ , called shared fitness  $s_f$ , is given by

$$s_f(i) = \frac{f(i)}{m(i)} \quad (11.24)$$

where  $f(i)$  is the original fitness and  $m(i)$  is called the niche count. For an individual  $i$ , the quantity  $m(i)$  is calculated by summing up the sharing function values contributed by all  $N$  individuals of the population,

$$m(i) = \sum_{j=1}^N \text{sh}(d_{ij}) \quad (11.25)$$

where  $d_{ij}$  is the distance between two individuals  $i$  and  $j$  and

$$\text{sh}(d_{ij}) = \begin{cases} 1 - \left(\frac{d_{ij}}{\sigma_s}\right)^\alpha & \text{if } d_{ij} < \sigma_s \\ 0 & \text{otherwise} \end{cases} \quad (11.26)$$

The quantities  $\alpha$  and  $\sigma_s$  are constants. A difficulty of this method is in choosing the value of  $\sigma_s$ . This requires prior knowledge about the number of peaks in the solution space. In our economical application as well as in many realistic problems, this information is not readily available.

A method is proposed in Yin and Gernay (1993) based on a different sharing scheme and using an adaptive cluster methodology. The authors show that this

Accepted for publication



method is effective in revealing unknown multimodal function structures and is able to maintain sub-population diversity. This method establishes analogies between clusters and niches in the following way. The GA population is divided by the adaptive MacQueen's KMEAN clustering algorithm, in  $K$  clusters of individuals that correspond to  $K$  niches. The shared fitness calculation is the same as in the classical sharing method, but the niche count  $m(i)$  is no longer associated with  $\sigma_s$ . In this case, the number of individuals within the cluster to which the individual  $i$  belongs plays a central role in the niche count calculation. As the number of clusters is associated with the number of niches (peaks), the individuals are put into a single partition of  $K$  clusters, where  $K$  is not fixed a priori but is determined by the algorithm itself. Therefore no a priori knowledge about the numbers of peaks of the fitness function is required as in the classical sharing method. The niche count  $m(i)$  is computed as

$$m(i) = N_c - N_c * \left( \frac{d_{ic}}{2 D_{max}} \right)^\alpha \quad x_i \in C_c \quad (11.27)$$

where  $N_c$  is the number of individuals in the cluster  $c$ ,  $\alpha$  is a constant, and  $d_{ic}$  is the distance between the individual  $i$  and the centroid of its niche. The algorithm requires a distance metric in order to compute the distance between two clusters and the distance between one individual and one cluster. Two clusters are merged if the distance between their centroids is smaller than a threshold parameter  $D_{min}$ . Moreover, when an individual is further away than a maximum distance  $D_{max}$  from all existing cluster centroids, a new cluster is formed with this individual as a member. The efficiency of the algorithm is improved by sorting the population in descending order according to the individual's fitness before the application of the clustering.

Such genetic algorithm with sharing and clustering has been applied to standard multimodal and continuous fitness functions by Yin and Gernay (1993) and Chopard *et al.* (1995) with promising results. One example of a more complex application was the determination of the optimum parameters of the *business* time scale,<sup>9</sup> which is used for analyzing price history and computing indicators. In this example, the optimization is quite difficult because we have to optimize simultaneously 17 parameters and the function is nonlinear in some of its parameters. To solve such a problem, it was necessary to add a normalization of the parameter space in the genetic algorithm—that is, each parameter is only allowed to vary in the range  $[0,1]$ . In simple problems, the two clustering parameters are generally set to  $D_{min} = 0.05$  and  $D_{max} = 0.15$ . But here, because of the large dimensionality of the parameter space, the value of the clustering parameters  $D_{min}$  and  $D_{max}$  must be much larger. In this case, the two parameters are multiplied by  $\sqrt{n}$  where  $n$  is the number of parameters to be optimized. The results obtained with this genetic algorithm are very promising and the sharing and clustering approach clearly increases

<sup>9</sup> This is a time scale that contracts and expands time based on seasonal activity or volatility of the time series (see Chapter 6).

ACCEPTED FOR PUBLICATION

the speed of convergence compared to the simple genetic algorithm described in the previous section.

When applied to the indicator optimization problem, the genetic algorithm with sharing and clustering runs into difficulties. If the fitness landscape contains too many sharp peaks of high fitness, all the selected clusters concentrate around these peaks and the genetic algorithm is unable to find robust solutions. In the next section, we propose some modifications to the genetic algorithm to detect clusters in the parameter space that correspond to more general and robust solutions.

*Modified Sharing Function for Robust Optimizations* We need to find a new genetic algorithm that avoids the concentration of many individuals around sharp peaks of high fitness but detects broad regions of the parameter space containing a group of individuals with a high average fitness level and a small variance of the individual fitness values.

To solve this problem, we propose a new sharing function that penalizes clusters with large variance of the individual fitness values and also penalizes clusters with too many solutions concentrated inside too small a region. The distance metric considered here is the Euclidean distance computed in the real parameter space (phenotypic sharing). In the proposed sharing scheme, all the individuals that belong to a given cluster  $c$  will share the same fitness value,

$$s_f(i) = \bar{f}_c - \left( \frac{N_c}{N_{av}} + \frac{1-r_d}{r_d} \right) \sigma(f_c) \quad \forall x_i \in C_c \quad (11.28)$$

where  $N_c$  is the number of genes in the cluster  $c$ ,  $\bar{f}_c$  is the average fitness value, and the standard deviation of the individual fitness values  $\sigma(f_c)$  is defined by

$$\bar{f}_c = \frac{1}{N_c} \sum_{i=1}^{N_c} f(i) \quad \text{and} \quad \sigma(f_c) = \sqrt{\frac{1}{N_c-1} \sum_{i=1}^{N_c} (f(i) - \bar{f}_c)^2} \quad (11.29)$$

As the method is based on the distribution of gene fitness inside each cluster, we keep only the clusters that contain at least a minimum number of members. We use a minimum cluster size of two individuals. As we also need to keep enough clusters of reasonable size, we have to limit the size of the largest clusters. The term  $N_c/N_{av}$  in Equation 11.28 is used to control the number of genes inside each cluster. If  $N_c$  is smaller than the expected average number of genes inside each cluster  $N_{av}$ , the correction is reduced, otherwise it is increased. Here, the constant  $N_{av}$  is chosen such that the population size is divided by the (preconfigured) expected number of clusters.

The second term  $(1-r_d)/r_d$  in Equation 11.28 is used to penalize clusters with too high a concentration of genes around their centroid. The value  $r_d$  is defined as

$$r_d = \sqrt{\frac{1}{N_c} \sum_{i=1}^{N_c} \left( \frac{d_{ic}}{D_{max}} \right)} \quad (11.30)$$

where  $d_{ic}$  is the distance of the gene  $i$  to the centroid of the corresponding cluster  $c$ . Here the square root is used to avoid too large a correction for an average concentration of genes, as is often the case.

To keep the cluster's space as large as possible, we also have to minimize the overlap between different clusters. To reduce this overlap, the clustering parameter  $D_{\min}$  must be quite large and here we use  $D_{\min} = D_{\max}$ . In order to have a reasonable clustering parameter for large dimensionality of the parameter space, the values of the two clustering parameters  $D_{\min}$  and  $D_{\max}$  are multiplied by  $\sqrt{n}$  where  $n$  is the number of parameters to be optimized.

With this new sharing scheme, the selection pressure is no more specific to each individual, as in a standard GA, but is the same for all genes present in a given cluster. This allows us to get a selection mechanism that looks for subpopulations of solutions with an average high quality instead of the best individual solution. Of course, the overall convergence speed is slightly reduced.

The selection pressure toward robust solutions is still present through the adaptive cluster methodology that tends to create clusters around a group of good individuals and through the reproduction technique, which uses elitism and mating restriction inside each cluster. Moreover, to keep a larger variety in the population, all the individuals who do not really belong to any clusters (i.e., who are further than the maximum distance  $D_{\max}$  from all existing cluster centroids) will have an unmodified fitness value. During the reproduction phase these individuals will have no mating restriction and generally a slightly higher selection probability.

To speed up the full process, the result of each different gene is stored and not recomputed when this gene appears again in the next generations. Moreover, the information of all the previously computed solutions can be used at the end to assess the reasonableness of the optimum solution.

Eventually, the algorithm selects, for each cluster, the best solution that is not farther than the distance  $D_{\max}/2$  from the cluster centroid. The final solution is the solution selected for the cluster that has the maximum average fitness corrected by the variance—that is, for the maximum value of  $\overline{f_c} - \sigma(f_c)$ .

The success of this type of genetic algorithm is still quite sensitive to the quality of the fitness measure but also to the normalization of the parameter space (i.e., to the quality of the metric used in the cluster construction). If the parameters do not have all the same sensitivity, this should also be reflected in the clustering algorithm. That is why we introduce the possibility of modifying the normalization of the parameter space, but in many applications this is not enough and some parameter mapping functions are needed. These functions depend on the specific problem to solve.

### 11.5.2 Testing Procedures

Strict optimization and testing procedures are a necessary condition to obtain robust trading strategies. The three main phases in the development of new trading models are as follows:

ESTER HART

- The development and optimization of new trading strategies
- The historical performance tests to select the strategies from data that were not used to optimize the models
- Real-time tests to confirm the performance of the selected models

The amount of historical data available for both the development (optimization) and the testing of a new trading model is always of finite size. On one side, to obtain meaningful and robust optimization results, a data sample as large as possible is requested. On the other side, the same is true for the statistical tests of the performance of a new model. Of course, the same data cannot be used for both the optimization and for the test of a trading strategy. The available historical data must be split into a minimum of two different sample periods. One period, named the *in-sample*, is used for the optimization and the other one, named the *out-of-sample*, for the performance tests. Such splitting must never be modified during the optimization or the testing phase, otherwise the risk of overfitting the historical data becomes very large and the statistical tests on the model performance are unreliable.

Another problem to take into account with financial data is the long-term heteroskedasticity (i.e., the presence of clusters which correspond to periods where the average price volatility is higher and other ones where the average price volatility is lower). As many trading models can react quite differently according to the average volatility of the market prices, it is not very convenient if the two selected data sets for optimization and testing present significant differences in their statistical properties.

A rule that provides reasonable results is to use two-thirds of the historical data for the optimization and one-third for the tests. The first part of the optimization data must be kept for the indicator initialization. The size of this initialization period, also named the *build-up* period, depends on the type of the indicators. In the case of exponential moving averages, the size of the initialization must be approximately 12 times larger than the range of the slower moving average.

At the end of the optimization process, the performance tests are executed once. If these performance tests do not provide good results, then the new trading model must be rejected.

It is strongly recommended to avoid tiny modifications of the initial model until good performance tests are obtained, because such procedure implies that the out-of-sample data period is, in fact, used indirectly for the optimization process itself, and again opens the door to overfitting problems.

When a new model is selected and passes the historical performance tests, the final phase is to check it in real-time for a few months. These last tests, named *ex-ante* tests, are useful to confirm the historical performance of the model and to check its reaction to real-time data flow. At Olsen & Associates (O&A), only the models that pass with success, both the historical test and the real-time *ex-ante* period, are used for real trading.

## 11.6 STATISTICAL STUDY OF A TRADING MODEL

### 11.6.1 Heterogeneous Real-Time Trading Strategies

The idea of this section is to use some trading models developed at Olsen & Associates as a tool to study the market structure (work presented in Dacorogna *et al.*, 1995). These models act like filters that concentrate on typical price movements and give us information about the market itself. The hypothesis of a heterogeneous market leads to three conjectures:

1. In a heterogeneous market, no particular trading strategy is systematically better than all the others. Excess return can be gained for different trading profiles, so various ways of assessing the risk and return of trading models are needed.
2. The different geographical components of the FX market have different business hours according to different time zones and, on the assumption of the heterogeneous market hypothesis, different strategies. Therefore, there are disruptions in the market behaviors from one geographical component to the next. Trading models that do not explicitly analyze the geographical components can avoid these disruptions only by restricting their active hours to the normal business hours of one geographical market. For such models, trading 24 hr a day does not pay.
3. The most profitable models actively trade when many agents are active in the market (liquid periods) and do not trade at other times of the day and on weekends. The heterogeneous market hypothesis attributes the profitability of trading models to the simultaneous presence of heterogeneous agents, whereas the classical efficient market hypothesis relates this profitability to inefficiencies. (This would imply that the illiquid periods of the market are the most favorable for excess returns.) If our conjecture is right, the optimal daily trading time interval should depend on the traded FX rate rather than the model type. Trading will be most profitable when the main markets for a particular rate are active.

Two trading models based on different algorithms are used in this study. The performance of these models is analyzed against changing market conditions, trading intervals, opening and closing times, and market holidays. The first trading model (RTT) is the one described in Section 11.4.1. Whereas the RTT model relies on one indicator with one time horizon, the second type of trading model (named here RTM) uses three different time horizons simultaneously to incorporate the views of three different market components. Like the RTT model, the RTM models have a profit objective of 3%, but the stop-loss value and profit objective are much smaller. The dealing frequencies of the RTM models are often higher than those of the RTT models, and they are also neutral more often.

The study presented here does not try to optimize the models in any way, so the distinction between in and out-of-sample is of little relevance. All the tests were conducted in a 7-year period from March 1986 to March 1993 for

ACCEPTED FOR PUBLICATION

**TABLE 11.3** Performance comparison between models.

Performance comparison between the O&A class RTT, RTM trading models, and a (benchmark) 20-day moving average model. The displayed performance measures are the annualized total return R, the risk-sensitive performance measure  $X_{eff}$ , the maximum drawdown D, the profit-loss ratio P/L and the dealing frequency F. These performance measures are explained in Sections 11.2.2 and 11.3.1.

FX rate	Model	R	$X_{eff}$	D	P/L	F
USD-DEM	MA(20)	5.5%	-0.9%	21.1%	0.57	1.0
	RTT	16.9%	11.2%	9.6%	0.41	1.7
	RTM	11.3%	8.6%	8.4%	0.68	2.0
USD-JPY	MA(20)	6.6%	0.6%	21.3%	0.53	0.9
	RTT	9.6%	4.2%	10.9%	0.59	1.5
	RTM	6.0%	3.5%	9.6%	0.45	1.9
GBP-USD	MA(20)	10.7%	5.5%	14.0%	0.58	1.0
	RTT	11.9%	7.1%	14.6%	0.40	1.6
	RTM	10.6%	8.2%	7.9%	0.66	2.1
USD-CHF	MA(20)	8.0%	0.9%	19.2%	0.59	1.1
	RTT	11.6%	6.1%	14.5%	0.55	1.3
	RTM	14.0%	10.1%	16.9%	0.65	1.9
USD-FRF	MA(20)	7.1%	4.0%	15.8%	0.56	1.0
	RTT	15.5%	11.2%	7.5%	0.75	1.1
	RTM	10.7%	8.6%	5.3%	0.60	2.1
USD-NLG	MA(20)	7.5%	3.3%	16.6%	0.55	1.0
	RTT	16.4%	10.9%	8.7%	0.50	1.7
	RTM	14.0%	11.2%	7.4%	0.69	2.1
USD-ITL	MA(20)	8.5%	1.7%	21.7%	0.57	1.0
	RTT	14.6%	7.2%	10.5%	0.42	1.6
	RTM	9.4%	6.1%	9.3%	0.65	1.9
DEM-JPY	MA(20)	1.4%	0.8%	4.9%	2.00	0.1
	RTT	10.9%	8.7%	6.5%	0.66	1.9
	RTM	10.1%	8.6%	5.9%	0.73	1.6
Average	MA(20)	7.7%	2.2%	18.5%	0.56	1.0
	RTT	13.4%	8.3%	10.4%	0.54	1.6
	RTM	10.8%	8.1%	8.8%	0.64	2.0

**TABLE 11.4** Performance comparison between markets.

The average risk-adjusted return  $X_{eff}$  for the different markets is shown as a percentage the average dealing frequency  $F$  is given in number of deals per week. The markets are listed in the order of their opening times in GMT.

Market	$X_{eff}$		Dealing frequency (F)	
	RTT	RTM	RTT	RTM
Tokyo	-0.8	1.6	1.3	1.4
Singapore	-0.4	2.3	1.4	1.4
Frankfurt	7.5	6.7	1.5	1.8
Vienna	8.3	7.8	1.5	2.0
Zurich	8.3	8.1	1.6	2.0
London	8.6	8.2	1.6	2.0
New York	6.3	6.7	1.5	1.9

USD-DEM, USD-JPY, GBP-USD, USD-CHF and DEM-JPY, and a  $6\frac{1}{2}$  year period from December 1986 to March 1993 for USD-FRF, USD-NLG, and USD-ITL.

Table 11.3 shows the comparative performance of the two types of models (RTT and RTM) together with the performance of a simple 20-day moving average model tested with the same high frequency data and the same environment. All models produce a significant profit even when transaction costs are fully accounted for. However, they differ both in the size of the average profit and in the risk of temporary losses. This was formulated as the first conjecture in the introduction. These results are a good illustration of the possibility of having diversified strategies that are all profitable but correspond to different risk profiles.

Realistic trading models should be configured for traders located in particular geographical locations. Our high-frequency data give us the flexibility of configuring different opening hours for different markets. In Table 11.3, the models were computed within the market constraints of Zurich. Now we want to show how the effective return varies if the market constraints are changed. Six other markets are tested: Frankfurt, London, New York, Singapore, Tokyo, and Vienna. Table 11.4 shows the different parameters related to the active times of these markets.

The same eight FX rates used for the performance comparison in Table 11.3 were tested here. In Table 11.4, we present the average of  $X_{eff}$  over the eight FX rates for the seven markets and the corresponding mean dealing frequency. The bad results for Tokyo and Singapore are not surprising because these markets are the least liquid. Good results in these markets are only obtained for USD-JPY and DEM-JPY. For the other five markets,  $X_{eff}$  generally does not vary much (within 1 to 2%), but it clearly peaks on the most active market (London) although the models were optimized for the Zurich market. This presents the first empirical evidence for our third conjecture that within the active times, the performance is not very sensitive to certain changing conditions.

TABLE 11.4



**TABLE 11.5** The best  $X_{eff}$  as a function of opening hours.

The best  $X_{eff}$ , in percent, as a function of the number of daily business hours and the opening and closing times in MET. The sixth column shows the  $X_{eff}$  reached when the models are allowed to trade 24 hr. The last column shows the hour that produces the best  $X_{eff}$  when only 1 hour per day is allowed for trading.

FX rate	Model	Best $X_{eff}$	Interval size	Daytime	$X_{eff}(24hr)$ 24hr trading	Best 1-hr trading
USD-DEM	RTT	9.7	12hr	7:00 – 19:00	6.7	16:30 – 17:30
	RTM	9.0	11hr	8:30 – 19:30	2.8	12:00 – 13:00
USD-JPY	RTT	4.3	9hr	8:30 – 17:30	-4.2	10:00 – 11:00
	RTM	7.6	9hr	3:00 – 12:00	0.5	17:00 – 18:00
GBP-USD	RTT	13.4	10hr	9:00 – 19:00	6.9	16:30 – 17:30
	RTM	9.0	8hr	6:00 – 14:00	5.8	13:00 – 14:00
USD-CHF	RTT	7.4	1hr	16:30 – 17:30	-5.8	16:30 – 17:30
	RTM	5.1	8hr	9:30 – 17:30	-0.6	18:30 – 19:30
USD-FRF	RTT	11.2	8hr	11:00 – 19:00	0.5	17:00 – 18:00
	RTM	9.4	9hr	8:00 – 17:00	1.6	16:00 – 17:00
USD-NLG	RTT	12.3	8hr	12:00 – 20:00	6.9	16:00 – 17:00
	RTM	8.8	10hr	8:30 – 18:30	4.6	15:30 – 16:30
USD-ITL	RTT	9.2	9hr	12:00 – 21:00	0.3	16:30 – 17:30
	RTM	6.6	8hr	11:00 – 19:00	1.3	13:00 – 14:00

At the beginning of this Section 11.6.1, we introduced two conjectures as subjects of research: it is not favorable to extend the dealing period to more than the normal business hours or even to 24-hr trading for our model types (conjecture 2); and the most profitable dealing periods should be the most active and liquid ones (conjecture 3). To test these conjectures, two main questions were asked: Is there an optimal daily business interval and do these optimal opening and closing hours differ for different rates? We present here the results of a study where we vary both the length and the starting point of the daily opening period. The two real-time model classes were tested with daily working intervals of 1, 8, 9, 10, 11, 12, 13, and 24 hr, shifting the opening time in 30-min steps from 0:00 to 24:00.

Table 11.5 shows the best  $X_{eff}$  values together with their corresponding working hours in Middle European Time (MET) for all rates and trading models used in this study. The models were optimized in-sample on  $9\frac{1}{2}$  hours from 8:00 to 17:30. Some first remarks can be made by looking at the results: shorter time intervals (8–10 hr) are generally preferred to longer ones (11–13 hr), thus confirming conjecture 2. There is not much profit in long working time intervals; these only tend to increase the number of bad deals because the indicators are more sensitive to



noise. Yet, because longer time intervals cover a larger period, the  $X_{eff}$  values of longer time intervals are more stable against changing opening and closing hours, that is, their variance is clearly smaller than that of shorter intervals. One exception is model RTT for USD-CHF where the 1-hr time interval is best, but the models still have significant peaks at the 9-hr interval (from 8:30 to 17:30, 4.9%).

Further evidence in support of conjecture 2 is given by the sixth column in Table 11.5, listing the  $X_{eff}$  values for 24-hr trading for comparison with the best  $X_{eff}$  values attained for shorter trading intervals. The  $X_{eff}(24\text{-hr})$  values are much lower than the best  $X_{eff}$  values for almost all rates and models. This failure of 24-hr trading can be interpreted as an insufficiency of the models to deal with short-term price movements, in particular the disruptive market behaviors arising when the main dealing activity shifts from one geographical location to another (with a different time zone). A 24-hr trading interval leads to a dealing frequency higher than that of a 12-hr interval. Contrary to the 24-hr trading interval, the best 1-hr intervals that coincide with the most active times of the day are seldom significantly worse than the rest of the intervals tested.

In Table 11.5, there are also indications that conjecture 3 is valid. The USD-JPY models show a strong tendency toward favoring opening hours early in the European morning (or closing times early in the afternoon), whereas GBP-USD and USD-ITL prefer opening times in the late morning. These results are in line with the time zones of the home markets of the currencies and must be related to market liquidity. For JPY, better results are obtained when its main market (Far East) is active, for GBP and ITL, when London (1-hr behind the Zurich market) is active (ITL is traded more in London than in Milan). The results of the RTT model for USD-NLG seem to contradict this conjecture, but it should be qualified. There is, in fact, another peak with an  $X_{eff}$  of 11.3% at a 12-hr trading interval from 7:30 to 19:30.

In conclusion, the systematic analysis of the influence of the trading hours on these models reveals some important facts. First of all, if we regard our model classes as representing medium-term components in the market, we see that it is not useful to stay active 24-hr a day. Without a much more sophisticated treatment of the intraday movements, it does not pay for a medium-term trader to be active all the time. Second, it shows that, contrary to assumptions based on the classical efficient market hypothesis, a trading model is profitable when its active hours correspond to the most active hours of one of the main geographical components of the market. It is essential that the models execute their deals when the market is most liquid. This fact is illustrated by three empirical findings:

- The maxima of performance are clustered around opening hours when the main markets are active.
- The best active times are shifted for certain currencies to accommodate their main markets (Japan for JPY, London for GBP).
- If the models are only allowed to trade for 1 hour, the best choice of this hour is usually around the peaks in the daily activity of the market.

2020 RELEASE UNDER E.O. 14176

The systematic variation of the business hours of the trading models again reveals the geographical structure of the FX market and its daily seasonality by the most profitable trading times being concentrated where the market is most liquid.

### 11.6.2 Price-Generation Processes and Trading Models

Instead of feeding the trading models with real data, we can use simulated data from different price-generation processes. The results indicate that the performance of the trading models with real FX data is much higher relative to the simulated price processes. This demonstrates that the trading models successfully exploit a certain predictability of returns that exists beyond the scope of the studied price-generation processes. The results also provide opportunity to compare different statistical price processes with each other. In the case of the RTT model described in Section 11.4.1, the out-of-sample test period is 7 years of high-frequency data on three major foreign exchange rates against the U.S. Dollar and one cross rate. From its launch in 1989 until the end of 1996, the model had not been reoptimized and was running on the original set of parameters estimated with data prior to 1989. This allows us a unique advantage that there is no socially determined coevolutionary relationship between our data set and the technical strategies used in implementing our specification tests.

The trading model yields positive annualized returns (net of transaction costs) in all cases. Performance is measured by the annualized return,  $X_{eff}$ ,  $R_{eff}$ , deal frequency and maximum drawdown. Their simulated probability distributions are calculated with the three traditional processes, the random walk, GARCH, and AR-GARCH, but also with an AR-HARCH. The null hypothesis of whether the real-time performances of the foreign exchange series are consistent with these traditional processes is tested under the probability distributions of the performance measures. As expected from the discussions of the previous chapters, the results from the real-time trading model are not consistent with the random walk, GARCH(1,1) and AR-GARCH(1,1) as the data-generating processes. It is also the case with the AR-HARCH processes.

*Simulation Methodology* The distributions of the performance measures under various null processes are calculated by using a simulation methodology. In our trading model simulations, we use a 5-min interval sampling of the prices in order to keep the computation within manageable bounds. It is a good compromise between efficient computation and realistic behavior when compared to the real-time trading model results generated from all ticks. The main information used by a trading model to update its indicators is the *returns*. The return between two consecutive selected ticks at time  $t_{j-1}$  and  $t_j$  is defined as

$$r_j = x_j - x_{j-1}$$

and the corresponding elapsed  $\vartheta$ -time (described in Section 6.2) between these two ticks is

$$\Delta\theta_j = \theta_j - \theta_{j-1}$$

By construction, in the sampled time series, the average elapsed  $\vartheta$ -time between two ticks,  $\overline{\Delta\theta}$ , is nearly 5 min.

Multiple time series from a given theoretical price generation process need to be generated. To keep the impact of special events like the data holes in the model behavior, we decided to replace the different bid-ask price values but always keep the recorded time values. As the different ticks are not exactly regularly spaced, even in  $\vartheta$ -time, the average return corresponding to a 5-min interval needs to be calculated. This is calculated by rescaling the observed return values

$$r_j^* = r_j \left( \frac{\overline{\Delta\theta}}{\Delta\theta_j} \right)^{1/E}$$

where the exponent  $1/E$  is called the drift exponent and it is set to 0.5 under the random walk process.

To obtain meaningful results, a simulated time series should have the same average drift  $\alpha$  and average variance  $\sigma^2$  as the observed returns. This is done by generating returns,  $\hat{r}_j$ , corresponding to a 5-min interval in  $\vartheta$ -time. In the case of a random walk process, the returns  $\hat{r}_j$  are computed with

$$\hat{r}_j = \alpha + \epsilon_j$$

where  $\epsilon_j \sim N(0, \sigma^2)$ .<sup>10</sup> When the effective elapsed time between two ticks,  $\Delta\theta_j$ , is not exactly 5-min, we scale again the generated return using the same scaling formula

$$r'_j = \hat{r}_j \left( \frac{\Delta\theta_j}{\overline{\Delta\theta}} \right)^{1/2}$$

where  $\overline{\Delta\theta}$  is 5 min. If there is a data hole, the sum of the generated return  $\hat{r}_i$  is computed until the sum of the added 5-min intervals is larger than the size of the data hole measured in  $\vartheta$ -time. The sum of the returns is scaled with the same technique as individual returns.

The simulated logarithmic prices,  $x'_j$ , are computed by adding the generated returns  $r'_j$  to the first real logarithmic price value  $x_0$ . The bid-ask prices are computed by subtracting or adding half the average spread, that is,

$$p'_{ask,j} = \exp\left(x'_j + \frac{\bar{s}}{2}\right)$$

and

$$p'_{bid,j} = \exp\left(x'_j - \frac{\bar{s}}{2}\right)$$

<sup>10</sup> In the simulations,  $\epsilon$  is specified to be normally distributed. We also explored bootstrapping the residuals of the studied models. The main findings of the study remain unchanged between these two approaches.

The parameters and the normalized residuals of the GARCH(1,1) process are estimated using the maximum likelihood procedure presented in Chapter 8. The simulated returns are generated from the simulated normalized residuals and the estimated parameters. The estimated parameters of the AR(p)-GARCH(1,1) processes together with the simulated residuals are used to generate the simulated returns for this process. As before, half of the average spread is subtracted (added) from the simulated price process to obtain the simulated bid (ask) prices.

For each replication we start by generating the simulated data a year before the model is tested. This year is 1989 and it is used to create the history dependency in returns and to initialize the different trading model indicators.

*Empirical Results* The simulated data are the 5-min  $\vartheta$ -time series,<sup>11</sup> from January 1, 1990, to December 31, 1996, for three major foreign exchange rates, USD-DEM, USD-CHF (Swiss Franc), USD-FRF (French Franc), and the most liquid cross-rate DEM-JPY (Deutsche Mark – Japanese Yen). From Chapter 6, we know that high-frequency data inherits intraday seasonalities and require deseasonalization. We use for this study the deseasonalization methodology presented in Chapter 6. Our data set contains 671,040 observations per currency. The simulations for each currency and process are done for 1000 replications.

Before discussing the details of different studies, we present in Table 11.6 results that substantiate the claims made at the beginning of this section. We give a summary of the  $p$ -values of the main performance measures for the USD-DEM, USD-CHF, USD-FRF, and DEM-JPY. The  $p$ -value<sup>12</sup> represents the fraction of simulations generating a performance measure larger than the original.

The methodology of this study places a historical realization in the simulated distribution of the performance measure under the assumed process and calculates its one-sided  $p$ -value.<sup>13</sup> This indicates whether the historical realization is likely to be generated from this particular distribution or not. More important, it indicates

<sup>11</sup> The real-time system uses tick-by-tick data for its trading recommendations. The simulations in this study are carried out with 5-min data as it is computationally expensive to use the tick-by-tick data for the simulations. The historical performance of the currency pairs from the 5-min series are within a few tenths of a percent for all performance measures with the performance of the real-time trading models utilizing the tick-by-tick data. Therefore, there is no loss of generality from the use of the 5-min frequency for the simulations instead of the tick-by-tick feed.

<sup>12</sup> The  $p$ -value represents a decreasing index of the reliability of a result. The higher the  $p$ -value, the less we can believe that the observed relation between variables in the samples is a reliable indicator of the relation between the respective variables in the population. Specifically, the  $p$ -value represents the probability of error that is involved in accepting our observed result as valid, that is, as *representative of the population*. For example, a  $p$ -value of 0.05 indicates that there is a 5% probability that the relation between the variables found in our sample is *purely coincidental*. In other words, assuming that in the population there was no relation between those variables whatsoever, and by repeating the experiment, we could expect that in approximately every 20 replications of the experiment there would be one in which the relation between the variables in question would be equal or stronger than ours. In many areas of research, the  $p$ -value of 5% is treated as a *borderline acceptable* level.

<sup>13</sup>  $p$ -value calculations reported in this study are the *simulated*  $p$ -values obtained from the distribution of 1000 replications of a given performance measure. For brevity, we simply refer to it as  $p$ -value in the text.

**TABLE 11.6** *p*-value Comparisons.

*p*-value comparisons with random walk (RW), GARCH(1,1), and AR(4)-GARCH(1,1). The *p*-values are expressed in percentage. The definitions of the three performance measures are presented in Section 11.3.

Currency	RW	GARCH(1,1)	AR(4)-GARCH(1,1)
<i>Annual return</i>			
USD-DEM	0.3	0.4	0.1
USD-CHF	8.9	8.4	3.7
USD-FRF	1.2	0.9	0.3
DEM-JPY	2.1	1.2	0.5
<i>Xeffective</i>			
USD-DEM	0.0	0.1	0.1
USD-CHF	0.7	1.4	1.9
USD-FRF	0.2	0.1	0.2
DEM-JPY	0.2	0.4	0.1
<i>Reffective</i>			
USD-DEM	0.0	0.0	0.0
USD-CHF	0.6	0.9	2.3
USD-FRF	0.1	0.1	0.1
DEM-JPY	0.1	0.4	0.1

whether the historical performance is likely to occur in the future. A small *p*-value (less than 5%) indicates that the historical performance lies in the tail of the distribution and the studied performance distribution is not representative of the data-generating process, given that the trading model is a good one. If the process that generates the performance distribution is close to the data-generating process of the foreign exchange returns, the historical performance would lie within two standard deviations of the performance distribution, indicating that the studied process may be retained as representative of the data-generating process.

*Random Walk Process* The results for the random walk process for USD-DEM time series are reported in Table 11.7. The first and the second columns are the historical realization and the *p*-value of the corresponding performance measures. The remaining columns report the 5<sup>th</sup> and the 95<sup>th</sup> percentiles, mean, standard deviation, skewness, and the kurtosis of the simulations.

After the transaction costs, actual data with the USD-DEM, USD-CHF, USD-FRF and DEM-JPY yield an annualized total return of 9.63, 3.66, 8.20, and 6.43%, respectively. The USD-CHF has the weakest performance relative to the other three currencies. The  $X_{eff}$  and  $R_{eff}$  performance of the USD-DEM, USD-FRF, and DEM-JPY are all positive and range between 3 and 4%. For the USD-CHF, the  $X_{eff}$  and  $R_{eff}$  are -1.68 and -4.23%, reflecting the weakness of its performance.

**TABLE 11.7** Random Walk Simulations for USD-DEM.

The second column presents the performance of the trading model with the actual data. The results under columns  $p$ -value, percentile, mean, standard deviation, skewness, and kurtosis present the values of these statistics from 1000 replications with the random walk process computed every 5-min for a period from 1990 to 1996. The  $p$ -values are reported in percentage terms (e.g., 0.3 refers to 0.3%). The definitions of the performance measures are presented in Sections 11.2.2 and 11.3.

Description	Historical realization	$p$ -value (in %)	Percentile (5%, 95%)	Mean	St.Dev	Skew.	Kurt.
Annual return	9.63	0.3	-11.38, 4.03	-3.44	4.74	0.09	-0.13
Xeffective	3.78	0.0	-20.25, -4.14	-12.11	5.09	0.13	-0.23
Reffective	4.43	0.0	-26.42, -7.70	-16.80	5.90	0.03	-0.20
Max drawdown	11.02	100.0	25.26, 94.86	53.79	21.36	-0.71	0.21
Deal frequency	1.68	100.0	2.20, 2.71	2.46	0.16	-0.10	-0.19
Horizon:	7 days						
Xeffective	3.47	0.0	-19.65, -4.24	-11.83	4.76	0.08	-0.15
Reffective	1.80	0.0	-24.14, -7.21	-15.51	5.20	0.05	-0.15
Horizon:	29 days						
Xeffective	3.27	0.0	-20.21, -4.36	-12.10	4.95	0.07	-0.23
Reffective	2.16	0.0	-27.05, -8.07	-17.45	5.91	0.02	-0.28
Horizon:	117 days						
Xeffective	4.07	0.0	-20.85, -3.42	-12.21	5.44	0.10	-0.32
Reffective	5.10	0.0	-31.01, -6.53	-18.10	7.49	0.26	0.25
Horizon:	301 days						
Xeffective	4.62	0.0	-23.37, -2.42	-11.89	6.32	0.39	0.02
Reffective	6.83	0.0	-27.85, -3.25	-14.56	7.49	0.35	0.16

The  $p$ -values of the annualized return for the USD-DEM, USD-CHF, USD-FRF, and DEM-JPY are 0.3, 8.9, 1.2, and 2.1%, respectively. For the USD-DEM and USD-FRF, as reported in Table 11.6, the  $p$ -values are less than the 2% level and it is about 2% for the USD-CHF. In the case of the USD-CHF, the  $p$ -value for the annualized return is 8.9, which is well above the 5% level. As indicated in Section 11.3, the annualized return only utilizes two points of the equity curve leaving a large degrees of freedom to infinitely many paths that would be compatible with a given total return.  $X_{eff}$  and  $R_{eff}$  are more stringent performance measures, which utilize the entire equity curve in their calculations. The  $p$ -values of  $X_{eff}$  and  $R_{eff}$  are 0.0, 0.0% for USD-DEM, 0.7 and 0.6% for USD-CHF, 0.2 and 0.1% for USD-FRF, and 0.2 and 0.1% for DEM-JPY. The  $p$ -values for the  $X_{eff}$  and  $R_{eff}$  are all less than 1%, rejecting the null hypothesis that the random walk process is consistent with the data-generating process of exchange rate returns.

The maximum drawdowns for the USD-DEM, USD-CHF, USD-FRF and DEM-JPY are 11.02, 16.08, 11.36, and 12.03%. The mean maximum drawdowns

from the simulated random walk processes are 53.79, 63.68, 47.68, and 53.49 for the USD-DEM, USD-CHF, USD-FRF, and DEM-JPY, respectively. The mean of the simulated maximum drawdowns are three or four times larger than the actual maximum drawdowns. The deal frequencies are 1.68, 1.29, 1.05, and 2.14 per week for the four currency pairs from the actual data. The deal frequencies indicate that the RTT model trades on average no more than two trades per week although the data feed is at the 5-min frequency. The mean simulated deal frequencies are 2.46, 1.98, 1.65, and 3.08, which are significantly larger than the actual ones.

The values for the maximum drawdown and the deal frequency indicate that the random walk simulation will yield larger maximum drawdown and deal frequency values relative to the values of these statistics from the actual data. In other words, the random walk simulations deal more frequently and result in more volatile equity curves on average relative to the equity curve from the actual data. Correspondingly, the  $p$ -values indicate that the random walk process cannot be representative of the actual foreign exchange series under these two performance measures. The summary statistics of the simulated performance measures have negligible skewness and statistically insignificant excess kurtosis. This indicates that the distributions of the performance measures are symmetric and do not exhibit fat tails.

The simulation results with the random walk process demonstrate that the real-time trading model is a consistent model. In other words, a process with no mean and a homoskedastic variance should only perform to generate an average return that would match the mean transaction costs. This consistency property is an essential ingredient of a trading model and the real-time trading model passes this consistency test. The means of the simulations indicate that the distributions are correctly centered at the average transaction costs, which is expected under the random walk process. For instance, the mean simulated deal frequency of the USD-DEM is 2.46 deals per week or 127.92 ( $2.46 \times 52$ ) deals per year. The relative spread for the USD-DEM is 0.00025, which in turn indicates an average transaction cost of -3.20% per year. Given that the mean of the simulated annualized return is -3.44, we can conclude that the mean of the simulated annualized return distribution is centered around the mean transaction cost.

The behavior of the performance measures across 7-day, 29-day, 117-day and 301-day horizons is also investigated with  $X_{eff}$  and  $R_{eff}$ . The importance of the performance analysis at various horizons is that it permits a more detailed analysis of the equity curve at the predetermined points in time. These horizons correspond approximately to a week, a month, 4 months and a year's performance. The  $X_{eff}$  and  $R_{eff}$  values indicate that the RTT model performance improves over longer time horizons. This is in accordance with the low dealing frequency of the RTT model. In all horizons, the  $p$ -values for the  $X_{eff}$  and  $R_{eff}$  are less than a half a percent for USD-DEM, USD-FRF, and DEM-JPY. For USD-CHF, the  $p$ -values are less than 2.4% for all horizons. Overall, the multihorizon analysis indicates that the random walk process is not consistent with the data-generating process of the foreign exchange returns.

FOIA b 7 - Exempt



**TABLE 11.8** GARCH(1,1) parameter estimates.

The sample is 5-min returns from 1990–1996.

	USD-DEM	USD-CHF	USD-FRF	DEM-JPY
$\alpha_0$	4.95 (4.23)	0.11 (0.12)	9.38 (7.09)	2.97 (4.03)
$\alpha_1$	0.1111 (0.0005)	0.1032 (0.0007)	0.1572 (0.0007)	0.0910 (0.0005)
$\beta_1$	0.8622 (0.0007)	0.8578 (0.0009)	0.8137 (0.0009)	0.8988 (0.0006)
$LL$	6.45	6.17	6.29	6.34
$Q(12)$	4810	4201	4256	3089
$\hat{\epsilon}_{\sigma^2}$	1.04	1.03	1.07	1.05
$\hat{\epsilon}_{sk}$	-0.07	-0.03	-0.05	0.16
$\hat{\epsilon}_{ku}$	11.73	7.28	22.93	27.73

**GARCH(1,1) Process** A more realistic process for the foreign exchange returns is the GARCH(1,1) process, which allows for conditional heteroskedasticity. The GARCH(1,1) estimation results are presented in Table 11.8. The numbers in parentheses are the robust standard errors and the GARCH(1,1) parameters are statistically significant at the 5% level for all currency pairs. The Ljung-Box statistic is calculated up to 12 lags for the standardized residuals and it is distributed with  $\chi^2$  with 12 degrees of freedom. The Ljung-Box statistics indicate serial correlation for the USD-DEM. The variances of the normalized residuals are near one. There is no evidence of skewness but the excess kurtosis remains large for the residuals.

In Table 11.9, the simulation results with the GARCH(1,1) process are presented for the USD-DEM rate. Because GARCH(1,1) allows for conditional heteroskedasticity, it is expected that the simulated performance of the RTT model would yield higher  $p$ -values and retain the null hypothesis that GARCH(1,1) is consistent with the data-generating process of the foreign exchange returns. The results, however, indicate smaller  $p$ -values, which is in favor of a stronger rejection of this process relative to the random walk process.

One important reason for the rejection of the GARCH(1,1) process as well as the random walk model is that these are pure volatility processes without predictability of the direction of returns, which matters for trading models. Another reason is the aggregation property of the GARCH(1,1) process. The GARCH(1,1) process behaves more like a homoskedastic process as the frequency is reduced from high to low frequency. Because the RTT model trading frequency is less than two deals per week, the trading model does not pick up the 5-min level heteroskedastic structure at the weekly frequency. Rather, the heteroskedastic structure behaves as if it is measurement noise where the model takes positions, and this leads to the stronger rejection of the GARCH(1,1) as a candidate for the foreign exchange data-generating process.

11  
10  
9  
8  
7  
6  
5  
4  
3  
2  
1



TABLE 11.9 GARCH(1,1) simulations for USD-DEM.

Description	Historical realization	$p$ -value (in %)	Percentile (% , 95%)	Mean	St.Dev	Skew.	Kurt.
Annual return	9.63	0.4	-11.14, 5.12	-3.27	4.90	-0.08	-0.01
Xeffective	3.78	0.1	-20.40, -3.16	-11.88	5.18	-0.07	-0.11
Reffective	4.43	0.0	-26.60, -6.37	-16.50	6.10	-0.14	-0.05
Max drawdown	11.02	100.0	24.17, 93.96	53.33	21.50	-0.73	0.30
Deal frequency	1.68	100.0	2.14, 2.64	2.39	0.15	-0.02	-0.15
Horizon:	7 days						
Xeffective	3.47	0.2	-19.56, -3.49	-11.64	4.90	-0.06	-0.03
Reffective	1.80	0.0	-24.19, -6.58	-15.37	5.38	-0.08	-0.05
Horizon:	29 days						
Xeffective	3.27	0.2	-19.95, -3.29	-11.86	5.00	-0.12	-0.04
Reffective	2.16	0.1	-26.92, -6.75	-17.20	6.04	-0.20	-0.03
Horizon:	117 days						
Xeffective	4.07	0.1	-21.24, -2.77	-11.91	5.56	0.03	-0.28
Reffective	5.10	0.1	-30.17, -5.44	-17.57	7.60	0.29	0.31
Horizon:	301 days						
Xeffective	4.62	0.2	-22.73, -1.48	-11.73	6.42	0.28	0.16
Reffective	6.83	0.3	-27.64, -2.01	-14.28	7.72	0.26	0.38

In a GARCH process, the conditional heteroskedasticity is captured at the frequency that the data have been generated. As it is moved away from this frequency to lower frequencies, the heteroskedastic structure slowly dies away leaving itself to a more homogeneous structure in time. More elaborate processes, such as the multiple horizon ARCH models (as in the HARCH process of Müller *et al.* (1997a)), possess conditionally heteroskedastic structure at all frequencies in general. The existence of a multiple frequency heteroskedastic structure seems to be more in line with the heterogeneous structure of the foreign exchange markets.

Table 11.6 we presented a summary of the  $p$ -values of the annualized return for the USD-DEM, USD-CHF, USD-FRF and DEM-JPY. In the case of the GARCH(1,1) simulation, they are 0.4, 8.4, 0.9, and 1.2%, respectively. All four currency pairs except USD-CHF yield  $p$ -values, which are smaller than 1.3%. The  $X_{eff}$  and  $R_{eff}$  are 0.1 and 0.0% for USD-DEM, 1.4 and 0.9 percent for USD-CHF, 0.1 and 0.1% for USD-FRF, and 0.4 and 0.4 percent for DEM-JPY.

The historical maximum drawdown and deal frequency of the RTT model is smaller than those generated from the simulated data. The maximum drawdowns for the USD-DEM, USD-CHF, USD-FRF, and DEM-JPY are 11.02, 16.08, 11.36, and 12.03 for the four currencies. The mean simulated drawdowns are 53.33, 60.58, 46.00, and 48.77 for the four currencies. The mean simulated maximum drawdowns are three to four times larger than the historical ones. The historical deal frequencies are 1.68, 1.29, 1.05, and 2.14. The mean simulated deal frequencies are 2.39, 1.87, 1.59, and 2.66 for the four currencies. The differences between

the historical deal frequencies and the mean simulated deal frequencies remain large. Therefore, the examination of the GARCH(1,1) process with the maximum drawdown and the deal frequency indicates that the historical realizations of these two measures stay outside of the 5% level of simulated distributions of these two performance measures.

The mean simulated deal frequency for the USD-DEM is 2.39 trades per week. In annual terms, this is approximately 124.28 deals per year. The half spread for the USD-DEM series is about 0.00025 and this yields 3.11% when multiplied with the number of deals per year. The -3.11% return would be the annual transaction cost of the model. For the model to be profitable, it should yield more than 3.11% per year. Table 11.9 indicates that the RTT model generates an excess annual return of 9.63%, whereas the mean of the annualized return from the GARCH(1,1) process stays at the -3.27% level.

The multi-horizon examination of the equity curve with the  $X_{eff}$  and  $R_{eff}$  performance measures indicates that the GARCH(1,1) process as a candidate for the data generation mechanism is strongly rejected at all horizons from a 7-day horizon to a horizon as long as 301 days. The overall picture coming out of the test is not very different for the GARCH(1,1) than that of the random walk process.

*AR(4)-GARCH(1,1) Process* A further direction is to investigate whether a conditional mean dynamics with GARCH(1,1) innovations would be a more successful characterization of the dynamics of the high-frequency foreign exchange returns. The conditional means of the foreign exchange returns are estimated with four lags of these returns. The additional lags did not lead to substantial increases in the likelihood value.

The results of the AR(4)-GARCH(1,1) optimization are presented in Table 11.10. The numbers in parentheses are the robust standard errors and all four lags are statistically significant at the 5% level. The negative autocorrelation is large and highly significant for the first lag of the returns. This is consistent with the high-frequency behavior of the foreign exchange returns and is also observed in Dacorogna *et al.* (1993). The Ljung-Box statistics still indicate serial correlation in the normalized residuals. The variances of the normalized residuals are near one. There is no evidence of skewness but the excess kurtosis remains large for the residuals.

The  $p$ -values of the annualized returns are presented in Table 11.6. They are 0.1, 3.7, 0.3, and 0.5% for the USD-DEM, USD-CHF, USD-FRF, and DEM-JPY. The results indicate that the AR(4)-GARCH(1,1) process is also rejected under the RTT model as a representative data generating process of foreign exchange returns. Here again, a possible explanation of this failure is the relationship between the dealing frequency of the model and the frequency of the simulated data. The AR(4)-GARCH(1,1) process is generated at the 5-min frequency but the model dealing frequency is between one or two deals per week. Therefore, the model picks up the high-frequency serial correlation as noise and this serial correlation works against the process. This cannot be treated as a failure of the RTT model.

F O R E I G N E X C H A N G E

TABLE 11.10 AR(4)-GARCH(1,1) parameter estimates.

The sample is 5-min return from 1990–1996.  $\alpha_0$  values are  $10^{-9}$ . The numbers in parentheses are the standard errors. The standard errors of  $\alpha_0$  are  $10^{-11}$ .  $LL$  is the average log likelihood value.  $Q(12)$  refer to the Ljung-Box portmanteau test for serial correlation and it is distributed  $\chi^2$  with 12 degrees of freedom. The  $\chi^2_{0.05}(12)$  is 21.03.  $\hat{\epsilon}_{\sigma^2}$ ,  $\hat{\epsilon}_{sk}$  and  $\hat{\epsilon}_{ku}$  are the variance, skewness, and the excess kurtosis of the residuals.

	USD-DEM	USD-CHF	USD-FRF	DEM-JPY
$\alpha_0$	3.90 (3.40)	8.19 (9.03)	7.28 (5.80)	2.92 (3.93)
$\alpha_1$	0.099 (0.0005)	0.0874 (0.0006)	0.1349 (0.0007)	0.088 (0.0005)
$\beta_1$	0.8796 (0.0006)	0.8833 (0.0007)	0.8411 (0.0008)	0.9008 (0.0006)
$\gamma_1$	-0.176 (0.001)	-0.208 (0.001)	-0.200 (0.002)	-0.130 (0.002)
$\gamma_2$	-0.011 (0.001)	-0.031 (0.002)	-0.025 (0.002)	-0.090 (0.002)
$\gamma_3$	0.003 (0.001)	-0.001 (0.002)	-0.005 (0.002)	-0.005 (0.002)
$\gamma_4$	-0.004(0.001)	-0.002 (0.001)	-0.008 (0.002)	-0.010 (0.002)
$\bar{LL}$	6.46	6.19	6.30	6.35
$Q(12)$	623	531	492	374
$\hat{\epsilon}_{\sigma^2}$	1.04	1.03	1.07	1.05
$\hat{\epsilon}_{sk}$	-0.07	-0.04	-0.05	0.15
$\hat{\epsilon}_{ku}$	12.29	7.86	21.84	27.98

Rather, this strong rejection is evidence of the failure of the temporal aggregation properties of the AR(4)-GARCH(1,1) process at lower frequencies.

The rejection of the AR(4)-GARCH(1,1) process with the  $X_{eff}$  and  $R_{eff}$  is even stronger and very much in line with the results for the random walk and the GARCH(1,1). The  $p$ -values of the  $X_{eff}$  and  $R_{eff}$  are 0.1, 0.0 percent for USD-DEM, 1.9, 2.3% for USD-CHF, 0.2, 0.1% for USD-FRF, and 0.1, 0.1% for DEM-JPY. The  $p$ -values remain low at all horizons for the  $X_{eff}$  and  $R_{eff}$ . The  $p$ -values of the maximum drawdown and the deal frequency also indicate that in almost all replications the AR(4)-GARCH(1,1) generates higher maximum drawdowns and deal frequencies.

**Conclusions** This extensive analysis of real-time trading models with high-frequency data suggests two main conclusions. First, technical trading models can generate excess returns, which are explained neither by traditional theoretical processes nor by luck. Second, the foreign exchange rates contain conditional mean dynamics that are neither present in the random walk nor GARCH(1,1), and AR-GARCH(1,1) processes.

The dealing frequency of the model is approximately between one and two per week although the data feed is at the 5-min frequency. Because the model's trading frequency is less than two deals per week, it does not pick up the 5-min level heteroskedastic structure at the weekly frequency. Overall, the results presented in this section have a general message to the standard paradigm in econometrics. It is

TABLE 11.11 AR(4)-GARCH(1,1) simulations for USD-DEM.

Description	Historical realization	<i>p</i> -value (in %)	Percentile (5%, 95%)	Mean	St.Dev	Skew.	Kurt.
Annual return	9.63	0.1	-10.46, 3.13	-3.68	4.13	-0.01	-0.16
Xeffective	3.78	0.1	-16.72, -3.16	-9.95	4.27	-0.02	-0.18
Reffective	4.43	0.0	-21.37, -5.28	-13.37	4.93	-0.07	-0.15
Max drawdown	11.02	100.0	21.73, 84.55	49.07	19.16	-0.59	0.03
Deal frequency	1.68	100.0	1.89, 2.35	2.12	0.14	-0.04	-0.26
Horizon:	7 days						
Xeffective	3.47	0.1	-16.53, -2.86	-9.72	4.13	0.01	-0.18
Reffective	1.80	0.2	-19.63, -4.95	-12.33	4.45	-0.01	-0.17
Horizon:	29 days						
Xeffective	3.27	0.1	-16.90, -3.24	-9.94	4.21	0.00	-0.11
Reffective	2.16	0.1	-21.54, -5.87	-13.67	4.87	-0.02	-0.02
Horizon:	117 days						
Xeffective	4.07	0.1	-17.37, -2.83	-9.97	4.50	0.00	-0.26
Reffective	5.10	0.0	-24.10, -4.75	-14.17	5.98	0.23	0.21
Horizon:	301 days						
Xeffective	4.62	0.1	-18.19, -2.01	-9.83	4.95	0.19	0.16
Reffective	6.83	0.1	-22.90, -2.91	-12.15	6.15	0.34	0.53

not sufficient to develop sophisticated statistical processes and choose an arbitrary data frequency (e.g., 1 week, 1 month, annual) claiming afterward that this particular process does a “good job” of capturing the dynamics of the data-generating process. In financial markets, the data generating process is a complex network of layers where each layer corresponds to a particular frequency. A successful characterization of such data generating processes should be estimated with models whose parameters are functions of *intra* and *inter*-frequency dynamics. In other fields, such as in signal processing, paradigms of this sort are already in place. Our understanding of financial markets would be increased with the incorporation of such paradigms into financial econometrics. Our trading model, within this perspective, helps us to observe this subtle structure as a diagnostic tool.

## 11.7 TRADING MODEL PORTFOLIOS

In the previous sections we have described what trading models are and how we can optimize and test them. In this section we will briefly study the combination of different trading model strategies into portfolios and discuss the particular case of currency risk hedging.

Any trading strategy is based on some specific indicators and decision rules and then will perform better in some market conditions. To reduce the risk implied by the use of such trading models, it is common to combine various trading strategies, which provide different trading signals for the same asset, in a

portfolio of models. As these models generally do not have the same clustering of good and bad periods, the overall risk is then reduced. But this is true only if the composition of the trading model portfolio is not changed too often. Dynamic modifications of the trading model portfolio, which keep a reasonable risk profile, are very hard to obtain and at O&A we advocate choosing static portfolio strategies where the dynamic behavior is left to the trading models themselves.

The optimal trading model portfolio strategy depends on certain decisions of the investor such as the choice of the investment assets, frequency of changes, and limits of risk and exposure. One of the main problems is the selection of the trading models to be used in such a portfolio. It is easy to test many different combinations and to select the best one, but such a procedure can produce undesirable results. In fact, during this selection procedure, the risk of overfitting particular historical data may again occur from the back door. To overcome these types of problems, it is often desirable to use an equally weighted portfolio—that is, a portfolio where the same proportion of capital or credit limit is invested in each trading model. Another possibility is to select the optimal trading model portfolio using a robust optimization procedure, like in trading model optimization (Section 11.5.1), but we will not discuss such an approach here.

Table 11.12 compares the performance obtained for a trading model portfolio, which corresponds to an equally weighted portfolio strategy, to the performance of the individual trading models on the same period. The analysis period is from January 1993 to December 1997. During this period, all these trading models were running in the real-time O&A information system with no reoptimizations. On this table we observe very well that for the same annualized total return the risk of the portfolio is considerably lower than the average risk on the individual models. The maximum drawdown of the portfolio is about half of the average maximum drawdown of the models and the annualized  $X_{eff}$  is one of the largest. The variation of the total return of the portfolio over the years is plotted on Figure 11.4.

Portfolios of trading models can be used as dynamic investment strategies in many financial markets, but the complexity of the optimization of such portfolios based on a very large number of trading models (needed for a good diversification) would be extremely hard to control. As we observed in the previous sections, the optimization of the different trading strategies is in itself a complicated process that needs to be done at regular intervals to take into account the nonstationarity of the underlying time series.

In the case of foreign exchange, an interesting application of portfolio trading models is the dynamic hedging of currency risk. In this case, the number of models to optimize is reduced and it is reasonable to consider a dynamic hedging strategy based on the trading recommendations. In the next section we will provide a brief description of this approach.

SCOTT W. O'NEILL

**TABLE 11.12** Portfolio performance of O&A trading models.


Performance comparison between 10 O&A class RTT, RTM trading models, and an equally weighted portfolio of the same models. The different performance measures displayed are the annualized total return R, the risk-sensitive performance measure  $X_{eff}$ , the maximum drawdown D, and the annualized Sharpe ratio S.

FX rate	Model	R	$X_{eff}$	D	S
USD-DEM	RTT	6.8%	2.3%	10.2%	0.73
	RTM	3.4%	1.3%	9.8%	0.51
DEM-JPY	RTT	1.7%	-2.6%	16.1%	0.19
	RTM	9.3%	6.4%	8.9%	1.32
USD-CHF	RTT	3.7%	-1.2%	10.9%	0.38
	RTM	2.5%	-0.3%	13.7%	0.36
USD-FRF	RTT	9.2%	4.7%	9.9%	1.03
	RTM	5.6%	2.5%	10.0%	0.71
GBP-DEM	RTT	5.6%	2.2%	14.3%	0.69
	RTM	5.0%	2.8%	8.1%	0.79
Average values		5.2%	1.8%	11.2%	0.71
Portfolio		5.2%	4.1%	5.5%	1.09

## 11.8 CURRENCY RISK HEDGING

Hedging problems arise whenever an investor, for example, a fund manager or a commercial organization, is holding foreign assets such as foreign securities over a period of time. The foreign assets are denominated in a foreign currency. The investor measures the performance of his/her investment in terms of the investor's home currency. The foreign assets have a degree of volatility in terms of their own currency. Due to the foreign exchange rate movements the volatility is, however, higher when expressed in terms of the investor's home currency. This implies additional risk. By additionally taking a short position in the foreign currency, this implicit foreign currency exposure can be compensated and the risk can be reduced; this is the basic idea of hedging. Whereas a constant short position is referred to as *static* hedging, this section deals with *dynamic* hedging where the foreign currency positions vary over time.

In this section, a strategy of hedging the foreign exchange (FX) risk associated with foreign investment is specified. As an innovative element of this strategy, real-time trading models are used. The whole strategy can then be called a dynamic overlay. To be successful, we need profitable trading models that are only weakly




**FIGURE 11.4** Total return of a portfolio of 10 O&A trading models over 5 years.

correlated or anticorrelated to the usual primary investments, because positive correlation would imply an increased risk.

An investor's risk/return decisions must be matched to the set of all possible investments (including dynamic allocation of capital to the foreign currencies), the *feasible set*. Figure 11.5 shows this feasible set as a shadowed region. The upper-left border of this set is termed the *efficient frontier*; those investment portfolios lying along this frontier deliver the maximum possible return for the minimum possible risk or the minimum risk for the set of best possible returns. The point at which an investor's indifference curve has a common tangent with the efficient frontier represents the best possible match between the investor's preferences and the possible investment portfolios. The right, dashed vertical line in Figure 11.5 indicates the expectation for the risk of a primary investment, which is left completely unhedged. A circle is drawn where this vertical line intersects with the horizontal line indicating the expected return of the primary investment.

The dashed vertical line to the left in Figure 11.5 indicates the reduction of risk achieved through an optimal static hedge of the primary investment, which usually implies short positions in all foreign currencies in which the foreign assets





**FIGURE 11.5** Set of feasible portfolios available to an investor when he/she implements a currency hedging strategy. The efficient frontier lies on the upper-left edge of the set (gray area), or along the darkened edge in the figure.

of the primary investment are denominated. In practice, this reduction in risk may be purchased at the price of slightly reduced returns due to transactions costs, thus the circle that intersects this risk line is lowered to slightly below the expected return of the primary investment in this example.<sup>14</sup> This optimal static hedge has succeeded in reducing risk. However, note that there may still be some distance between the static point and the efficient frontier of the feasible set, which is defined by the use of dynamic allocation to foreign currencies through trading models. The distance between the statically hedged portfolios and the efficient frontier marks the improvement that can be attained through a dynamic currency overlay using trading models.

The subject of currency hedging has been discussed in the literature for some time. In Froot (1993), full hedging is recommended for minimizing the risk due to *short-term* FX volatility. On the other hand, it is also shown that a lower amount of hedging or even no hedging is better for minimizing the risk of *long-term* value

<sup>14</sup> Whereas the transactions costs of a one-time, static FX transaction are minimal, the short positions in the foreign currencies may imply considerable costs of carry due to interest rate differential between the two currencies. These costs, which are sometimes also in favor of the investor, may lead to a distinct return difference between statically hedged and unhedged portfolios.



fluctuations of the investment. Depending on the time horizon, there is thus a certain range, a scope within which the hedging ratio can be chosen.

Levich and Thomas (1993a) go one step further. They hedge a position *dynamically* by varying the hedge ratios over time. They show that this is *profitable* as compared to no hedging or static hedging. In their most successful strategy, a “currency overlay” with many currencies involved, they change the hedge ratios by following simple “technical trading signals.”

In the overlay strategy described here, the allowed ranges of static hedging and the exposure due to dynamic hedging are limited.<sup>15</sup> Thus the main purpose of hedging, which is reducing the risk due to FX rate volatility, is maintained. On the other hand, we have a well-founded additional profit expectation, based on the profitability of the trading models and trading model portfolios.

### 11.8.1 The Hedging Ratio and the “Neutral Point”

Currency hedging means, for an investor who has bought a foreign asset such as equity of value  $s$ , holding a short position of size  $-s_h$  in the foreign currency in order to minimize the volatility of the value of his/her total position due to FX rate fluctuations. The hedging ratio  $h$  is defined as

$$h = \frac{s_h}{s} \quad (11.31)$$

The study by Froot (1993) shows that choosing the best  $h$ , the one that minimizes the total volatility, is not trivial and depends on the time horizon of the investor. For short-term investors, the best  $h$  is 1 or slightly less; for long-term investors, who hold their position over many years, the best choice of  $h$  is about 0.35. In Froot (1993), the hedging ratio  $h$  is assumed to be constant over time.

In our *dynamic* hedging approach (we follow here the method suggested in Müller *et al.*, 1997b), we want to vary  $h$  over time, following some real-time trading models to reach an additional profit or to reduce the risk of the primary investment expressed in the home currency. This requirement will, when used during optimization, automatically set limits to the type of hedging strategy to use. In the lack of a clear criterion like reducing the risk, some rules could be introduced to achieve desirable features like, for instance,  $h$  would not depart too much from the best value—that is, in most cases, between 0.35 and 1 according to Froot (1993). Each of the individual foreign currencies has its own hedging ratio  $h_i$ , but there is also the option of taking only one *global* hedging ratio  $h$  for all currencies. The following discussion applies to both the individual  $h_i$  and  $h$ .

In lack of an objective risk criterion, most investors set some limits  $h_{\min}$  and  $h_{\max}$  on the choice of the hedging ratio, often to satisfy some institutional constraints or to limit the risk if they have no other quantitative criterion to do this. A typical choice might be  $h_{\min} = 0$  (no hedging) and  $h_{\max} = 1$  (full hedging).

<sup>15</sup> Interested readers are referred to our internal paper where the methodology is described in detail, Müller *et al.* (1997b).

FOR OFFICIAL USE ONLY

In the middle between the extreme  $h$  values, we define the *neutral point*,  $h$ , of the hedging strategy to be

$$h_{\text{mid}} = \frac{h_{\text{min}} + h_{\text{max}}}{2} \quad (11.32)$$

and the total possible range  $\Delta h$  of dynamic hedging is

$$\Delta h = h_{\text{max}} - h_{\text{min}} \quad (11.33)$$

We have seen that the O&A real-time trading models vary their gearings between two limits called “gearing -1” and “gearing 1” as explained in Section 11.2.1. If the neutral point  $h_{\text{mid}}$  is chosen for static hedging and if there is only one foreign currency and one trading model (for the FX rate between that foreign currency and the home currency), these trading limits directly correspond to the limits  $h_{\text{min}}$  and  $h_{\text{max}}$ . However, the investor may decide to allow wider exposure limits for the dynamic positions than those of the static positions. The situation becomes more complicated in the presence of many currencies and many trading models.

### 11.8.2 Risk/Return of an Overlay with Static and Dynamic Positions

To solve the allocation problem, a basis of portfolio theory has to be applied to our particular overlay problem.

A portfolio  $\Pi$  can be written in terms of the sum of the primary investment ( $PI$ ), a static foreign exchange position placed in order to hedge foreign exchange risk ( $SH$ ) and a series of variable positions placed to dynamically hedge foreign exchange risk ( $DH$ )

$$\Pi = PI + SH + DH \quad (11.34)$$

In what follows, we refer to a portfolio diversified in  $i = 1$  to  $n$  currencies and dynamically hedged by  $j = 1$  to  $m$  trading models:

$$\Pi = \sum_{i=1}^n \alpha_i I_i + \sum_{i=1}^n h_i \alpha_i R_i^{FX} + \sum_{j=1}^m \omega_j R_j^{TM} \quad (11.35)$$

Expectation of portfolio returns, or changes in value, can then be written as  $E[\Delta \Pi]$ , or as a sum of the expectation values of its comprising parts:

$$E[\Delta \Pi] = E[\Delta I] + E[\Delta R^{FX}] + E[\Delta R^{TM}] \quad (11.36)$$

that is,

$$E[\Delta \Pi] = \sum_{i=1}^n \alpha_i E[\Delta I_i] + \sum_{i=1}^n h_i \alpha_i E[\Delta R_i^{FX}] + \sum_{j=1}^m \omega_j E[\Delta R_j^{TM}] \quad (11.37)$$

where

$\Delta I_i$  = the fractional return of the  $i^{\text{th}}$  underlying foreign investment component at a given time,  $t$ , and over a time horizon,  $\Delta t$ ;  $\Delta I_i \equiv \Delta I_i(\Delta t, t) \equiv (I_t - I_{t-\Delta t})/I_{t-\Delta t}$ , where  $I_i$  are in units of the home currency,

$\alpha_i$  = the amount of the portfolio allocated to the  $i^{\text{th}}$  currency in units of the home currency,

$h_i$  = the unitless static hedging ratio for the  $i^{\text{th}}$  foreign currency,

$\Delta R_i^{FX}$  = the portfolio returns due to fluctuations in the static hedge positions,

$\omega_i$  = the weights given to each trading model (in units of the home currency), and

$\Delta R_j^{TM}$  = the trading model return values.

The risk of the portfolio is characterized by the variance of the portfolio returns  $\sigma_{\Delta\Pi}^2$ :

$$\sigma_{\Delta\Pi}^2 = E \left[ (\Delta\Pi - E[\Delta\Pi])^2 \right] \quad (11.38)$$

### 11.8.3 Dynamic Hedging with Exposure Constraints

To compute the efficient frontier of the dynamic hedging strategy, we need to optimize the return given the risk or conversely optimize the risk given the return. Such an optimization is suitably done using the Lagrange multiplier technique. We can maximize the quantity<sup>16</sup>  $E[\Delta\Pi] - \lambda \sigma_{\Delta\Pi}^2$  or, more conveniently, minimize  $\sigma_{\Delta\Pi}^2 - \lambda E[\Delta\Pi]$ . The parameter  $\lambda$  is the Lagrange multiplier and can be varied from 0 (considering risk only) to high positive values (considering mainly return while keeping risk under some control) to get the whole efficient frontier. Each value of  $\lambda$  corresponds to a point on the efficient frontier. Let us call  $U$  the target function to be minimized,

$$U = \sigma_{\Delta\Pi}^2 - \lambda E[\Delta\Pi] \quad (11.39)$$

where  $E[\Delta\Pi]$  and  $\sigma_{\Delta\Pi}^2$  have been defined and expressed in Equations 11.37 and 11.38.

Some components of the total portfolio have free coefficients, which can be determined with the goal of minimizing the target function  $U$ : the hedging ratios  $h_i$  and the amounts of money (maximum exposures)  $\omega_j$  allocated to the trading models. These coefficients are generally subject to constraints:

1. The trading model sizes  $\omega_j$  must not be negative:  $\omega_j \geq 0$ . In fact, such negative  $\omega_j$  would mean doing the opposite of the trading recommendations. In these cases, this will lead to new transaction costs, which make the actual returns much worse than the formal results suggest. Therefore, we exclude negative trading model coefficients.

<sup>16</sup> This first quantity is very similar to  $X_{eff}$ , the risk-corrected return derived in Section 11.3.1. The optimization problem comes back to optimizing  $X_{eff}$  for different risk aversion constants.

2. The static hedging ratios  $h_i$  are limited to an allowed range between  $h_{\min,i}$  and  $h_{\max,i}$ . Generally there are regulations or risk management rules that do not allow the investors to take too extreme positions, therefore hedging ratios that are above  $h_{\max,i}$  or below  $h_{\min,i}$  are generally excluded.

The constraints make a direct solution of the allocation problem impossible. The minimization of  $U$  (a quadratic function of the coefficients) under linear constraints on the coefficients is a special case of *quadratic programming*. A technique developed by Markowitz (1959), based on the simplex method, can solve this problem; we follow Markowitz (1987). There is no local minimum of  $U$  in the space of the coefficients; this can be proven. The solution, once found, always represents the global minimum.

Currency-wise exposure limits may be less rigid than those on the hedging ratios. But exposure limits on foreign currencies are not simple, as they involve several assets together. The exposure constraints are linear in  $\omega_j$  and  $h_i$ . As the simpler constraints, they can be fully accounted for in the framework of the solution method presented by Markowitz (1987).

The whole efficient frontier can be obtained by minimizing  $U$  for different choices of the Lagrange multiplier  $\lambda$ . The left-hand side of the efficient frontier obviously starts at  $\lambda = 0$  where the return of the portfolio has no influence and only risk counts. The right-hand side limitation of the efficient frontier is less evident. For unconstrained dynamic strategies, the efficient frontier normally extends to infinity (when more and more money is allocated to trading models). A hedging strategy should obviously not deal with infinite risks. There should be a risk limit beyond which the strategies are considered unacceptable even if the formal exposure limits are not yet reached. But in the case of constraints, the efficient frontier generally has a genuine end on the right-hand side. Therefore, we use a "shooting" strategy for determining the  $\lambda$  value that approximately leads to the desired maximum risk at the right end of the efficient frontier. The method by Markowitz (1987) helps us to find this value.

#### 11.8.4 Concluding Remarks

This problem turns out to be rather difficult and requires many different inputs and programs to solve it in a practical way. The algorithm needs the following main ingredients:

- The user-defined investment goals such as primary investment and exposure limits.
- Time series of returns of all assets in the portfolio: FX rates, interest rates (or their differentials), and typical primary investments such as stock indices and bond indices of many currencies.
- Time series of returns of the trading models used for dynamic hedging (including transaction costs).
- The computation of mean returns and the covariance matrix of all relevant assets.

2025 RELEASE UNDER E.O. 14176

- A target function to be optimized, measuring risk and return, with a Lagrange multiplier that sets the balance between risk and return. This function depends on two types of parameters, which are static hedging ratios and trading model allocation sizes.
- A method to solve the quadratic programming problem with linear constraints (exposure limits).

In conclusion, we are able to solve the dynamic overlay problem under exposure constraints with a quite complex but well-understood algorithm. The result is the efficient frontier of feasible solutions, each with a particular risk/return profile. It is the choice of the investor to decide which portfolio on the efficient frontier he/she wants to follow.

FORWARDED

# 12

---

## TOWARD A THEORY OF HETEROGENEOUS MARKETS

At the end of this in-depth review of some of the techniques and models used with high-frequency data, there is clear evidence that price movements of foreign exchange rates and other financial assets for short to medium-term horizons are predictable to some extent. This is substantiated by a positive forecast quality and high real-time trading model returns (e.g. Dacorogna *et al.*, 1992; Pictet *et al.*, 1992; Gençay *et al.*, 2001c, 2002). More generally, financial returns of whatever asset substantially depart from the random walk model and are being predicted with some success by market participants.

Where does this sustained predictability originate? Are the real-time trading models, for instance, successful in capturing the inefficiencies of the foreign exchange (FX) market? Because this market is widely held to be the most efficient of the financial markets, does this success conflict with the theory of efficient markets, which precludes the ability to forecast and denies the existence of profitable trading models? Should we conclude from this evidence that markets are inefficient? We believe that we should rather adapt our theory of the financial market to the reality of the stylized facts and of markets that are very efficient in a newly defined way.

The motivation of this chapter is to explain why and how markets can be at the same time highly efficient and to some extent predictable. There are a number

of reasons for this that are all associated with market dynamics. We want to put in perspective the current theory of efficiency and suggest to move beyond it. This is one of the big challenges ahead in the theory of finance. Many researchers are working in this special field, such as the whole movement of “behavioral finance” around Robert Shiller,<sup>1</sup> or parts of the econophysics group and many others who see the need to find ways of moving from a rather static definition to a more dynamic one.

## 12.1 DEFINITION OF EFFICIENT MARKETS

In conventional economics, markets are assumed to be efficient if all available information is reflected in current market prices (Fama, 1970, 1991). Economists have embarked on weak, semi-strong, and strong-form efficiency tests. The weak-form tests investigate whether market prices actually reflect all available information. The semi-strong tests are based on so-called event studies, where the degree of market reaction to “news announcements” is analyzed. The strong-form tests, finally, analyze whether specific investors or groups have private information from which to take advantage. By and large, most studies conclude that the major financial markets are efficient and that all information is reflected in current prices. However, the conclusions of such studies have been bogged down by methodological questions: in particular, whether any observed departures from market efficiency are due to any genuine market inefficiency or whether a deficiency of the market pricing model is being used as a yardstick to compare actual with theoretical prices.

The inference that in an efficient market no excess return can be generated with trading models is based on the assumption that all investors act according to the rational expectation model (Shiller, 1989; Fama, 1970). If this assumption is wrong, the conclusion that forecasting is impossible is also questionable. The assumption of rational expectations has been called into question on various platforms and the idea of heterogeneous expectations has become of increasing interest to specialists. Shiller (1989), for example, argues that most participants in the stock market are not “smart investors” (following the rational expectation model) but rather follow trends and fashions. The modeling of “noise trader” has become a central subject of research in market microstructure models. On the FX market, there is much investigation of “speculative bubbles” and the influence of technical analysis on the dealer’s strategy (see, for example, Frankel and Froot, 1990). Some attention has also been caught by the possibility of time-varying expectations, which better reflect to our view of the market (Bekaert and Hodrick, 1992). Variation over time in expected returns poses a challenge for asset pricing theory because it requires an explicit dynamic theory in contrast to the traditional static capital asset pricing model (CAPM).

<sup>1</sup> See, for instance, Shiller (2000) where the author claims that the market agents are essentially acting irrationally.

In summary, the conclusion that financial asset prices are not predictable is based on three assumptions: market prices reflect all the information available, news and events that hit the market are normally distributed, and the market is composed of homogeneous agents. The two first assumptions are reasonable starting points for the definition. The third assumption poses a real problem. It is clear that all market agents have in fact bounded rationality. They cannot be omniscient and do not all enjoy the same freedom of action and access to the markets. Recent works by Kurz (1994) and Gouree and Hommes (2000) present new theoretical models to tackle this problem. Introducing the heterogeneity of agents can give rise to very interesting nonlinear effects in the models. They show that many of the price fluctuations can be explained by endogenous effects. Similar conclusions are reached by Farmer and Lo (1999) in their discussion of market efficiency. They base their analysis on a comparison with the evolution of ecological systems. Farmer (1998) develops a market model inspired by ecological systems that contains agents with various trading strategies.

## 12.2 DYNAMIC MARKETS AND RELATIVISTIC EFFECTS

We just saw that conventional economics makes its inferences on efficient markets on the basis of a model in which economic agents are entities that act according to the rational expectation strategy. Any differences in planning horizons, frequency of trading, or institutional constraints are neglected. However, there is substantial empirical evidence that investors have heterogeneous expectations, as noted in Müller *et al.* (1993a) and Müller *et al.* (1997a). Surveys on the forecasts of participants in the FX market reflect the wide dispersion of expectation at any point in time. The huge volume of FX trading is another indication reinforcing this idea because it takes differences in expectation among market participants to explain why they trade.<sup>2</sup> In Chapter 7, we presented the heterogeneous market hypothesis; at the end of this book the need for such a view becomes clear. It is the most elegant way to reconcile market efficiency with the stylized facts. Lux and Marchesi (1999) have developed simulation models of financial markets that include agents with different strategies (fundamentalists and chartists). They were able to show (Lux and Marchesi, 2000) that this model can reproduce most of the empirical regularities (fat tails, long memory, and scaling law) even though they use normally distributed news in their simulations.

The theoretical work on financial markets with heterogeneous agents has also gained momentum in the literature. Among this literature, Brock (1993), Brock and Kleidon (1992), Brock and LeBaron (1996), Brock and Hommes (1997), and Hommes (2000) investigate the underlying source for the structural heterogeneity of financial markets. Brock (1993) studies the interacting particles system theory to build structural asset pricing models. Brock and Hommes (1997) build a general theory of expectation formation, which nests rational expectations in an

<sup>2</sup> Over \$1500 billion US is traded every day in the different centers like Tokyo, London, and New York according to a survey taken every 3 years by the Bank for International Settlements (1999).



econometrically tractable system. Brock and Kleidon (1992) show how bid-ask spreads fluctuate over the day by firm size categories as a measure of “thickness” of the market. Brock and LeBaron (1996) stress not only the standard asymmetric information theory in matching key stylized facts, but also the importance of the role of multiple time scales. Hommes (2000) provides a review of recent work on heterogeneous agent financial theory.

There are many ways to describe heterogeneous expectations. We believe that the most promising approach is to differentiate the expectations according to their time dimension because we consider the *different time scales* of the market participants the key characteristic of the market. Some are short-term traders, others have long-term horizons with market makers at the short-term end of the scale and central banks at the long-term end. Contrary to the usual assumption, there is no privileged time scale in the market. The interaction of components with different time scales gives rise to characteristically relativistic effects<sup>3</sup> such as certain properties of volatility clusters, trend persistence, lag between interest rate adjustment, and FX rate adjustment. The latter is a good example of what conventional theory considers an inefficiency whereas we see it as an effect arising from the different time scales involved in the market. To take advantage of the lag in adjustment between interest rate and exchange rate moves, an investor needs to tie up his/her money for months or even years. This is a very long time for an FX trader. Some investors will thus tend to ignore these profit opportunities whereas others invest in them, as is testified by the development of managed currency funds based on this property. The combination of all of these effects ultimately enables the construction of successful forecasting and trading models.

In long time intervals, market price changes are “flatter” and have fewer relevant movements (trend changes) than in short-term intervals. The higher the resolution and the smaller the intervals, the larger the number of relevant price movements. The long and the short-term traders thus have different trading opportunities: the shorter the trading horizon, the greater the opportunity set. A market participant’s response to outside events should always be viewed as relative to one’s intrinsic opportunity set. A short-term trader does not react in the same way as a long-term trader. Economic decision makers, such as traders, treasurers, and central bankers, interpret the same information differently. The variation in perspective has the effect that specific price movements cannot lead to a uniform reaction; rather, they result in individual reactions of different components. In turn, these reactions give rise to secondary reactions, with the different components reacting to their respective initial response. Watching the intraday price movements, one clearly sees the sequences of secondary reactions triggered by the initial events. See, for example, Goodhart and Curcio (1991); Almeida *et al.* (1998) on news effect on the FX market or Franke and Hess (1997) on the

<sup>3</sup> We use here the term “relativistic” to express the dynamic interaction between different market components relative to each other rather than relative to the news that has impacted the market. These effects are sometimes called endogenous effects in the literature, Kurz (1994).

Deutsche Termin-Börse. The existence of different trading strategies in the market was also put forward in Chapter 7 to explain the HARCH effect of asymmetry in the information flow at different frequencies. LeBaron (2000) shows that introducing agents with different time horizons in his market model gives rise to heteroskedasticity effects in the resulting price volatility.

The delay with which the secondary reactions unfold is called the *relaxation time*. If diverse components with different time scales interact in the market, there is typically a mixture of long and short relaxation times following the impact of outside events. If different relaxation times are combined, the resulting autocorrelation decays hyperbolically or almost hyperbolically. This is a natural explanation of the long memory effects detected in financial markets. Dacorogna *et al.* (1993) studied the autocorrelation function for short-term absolute returns, confirmed the hyperbolic decay, and revealed that volatility clusters tend to have a longer memory than assumed by other studies of the subject. We saw in Chapter 7 that many studies confirm this effect.

There is yet another phenomenon, which originates from the fact that financial markets are spread worldwide. Economic and political news and trading activity are not stationary. They have a clear-cut pattern of moving around the world in a 24-hr cycle. The price data of foreign exchange rates reflect this in terms of a 24-hr seasonality in market volatility, Müller *et al.* (1990). This seasonality can be accounted for by introducing a *business time scale* such as in Dacorogna *et al.* (1993). The 24-hr cycle implies that market reactions to an event cannot be simultaneous and that there are distinct relaxation times following the event. Geographical components related to the business hours of the different trading centers must be added to the time components. The interaction of geographical components leads to behaviors such as the “heat wave” effect proposed by Engle *et al.* (1990).

### 12.3 IMPACT OF THE NEW TECHNOLOGY

The realization that there is value in the data to define an investment strategy has brought to life many new firms that specialize in modeling financial markets and in providing trading advices on the basis of technical models. The question is, of course, will the impact of the new technology be a passing phenomenon or will it have a long-term effect? As the relativistic phenomenon arises from the interaction of components with different time scales, it will remain appropriate as long as heterogeneous expectations continue to exist in the market. The interaction process may become more complex, but it cannot disappear.

News technologies enable users to identify additional trading opportunities to increase their profits. This quickens their pace of trading and contributes to higher market volume and liquidity. The improved liquidity lowers the spreads between bid and ask prices. Lower spreads imply lower transaction costs, which in turn increase the opportunity horizon for profitable trading. The new technology

FOR THE COURT

introduces a shift in perspective, with components starting to focus on more numerous short-term time intervals.

As components become increasingly short-term in their focus, the spectrum of short-term components increases. This has the effect that relative differences among components become more significant and the relativistic effects more pronounced. Contrary to accepted notions, which assume that sufficient buying power can “trade away” any phenomenon, the increased buying power will have the overall effect of enhancing the relativistic effects. Thus the very basis of our ability to forecast and build profitable trading models will be enhanced. This statement must be qualified in the sense that the reaction patterns will become increasingly diversified, and therefore more complex, and the speed of adjustment will increase requiring more and more sophisticated models.

#### 12.4 ZERO-SUM GAME OR PERPETUUM MOBILE?

Conventional thought has it that financial markets must be a zero-sum game. This is true if we take a static view. In reality, the financial markets are dynamic and they are highly complex.

Markets are a platform for components to take advantage of the diversity of interests. They are able to match their opposing objectives when one component buys and another component sells. The lower the friction, the easier a counterpart for a particular transaction is found and the larger, therefore, is the particular component's opportunity set. By being able to go ahead with a particular transaction, the flexibility of the respective components is increased and their profit potential improved.

The new technology fosters the ability of the market to provide an environment for the generation of wealth. As explained, interaction within the market gives rise to relativistic effects and relaxation times. To the extent that these relativistic effects are understood and incorporated into forecasting and trading model technology, market participants have the opportunity to generate additional profit or limit their losses. In our terminology, the profit that is generated is energy extracted from the market. Improved efficacy of component interaction generates additional energy and reduces the friction associated with buying and selling within the market. The process may be compared to the search for more efficient engines in the automobile industry where everybody gains from it in the long term.

Have we achieved a *perpetuum mobile*? The answer is clearly no. Like any other technological innovation, the new technology does not generate energy from nothing, but it does take advantage of the energy potential existing in the financial markets. By offering a service to the economic agents, financial markets are not closed systems but do get a permanent input of money. This makes them highly open systems in terms of energy. Besides, a lot of resources have been put into the new technology in the form of extensive research, development work, and hardware to treat the information. Numerous studies have shown that simple trading rules do not work in efficient markets. Only elaborated treatment of the

RECEIVED

data allows the identification of profitable trading rules. This treatment is not free, it has a price. Moreover, as the relativistic reaction patterns become increasingly diversified, research and development efforts will have to increase in the future to keep up with the ever-changing nonlinear patterns.

## 12.5 DISCUSSION OF THE CONVENTIONAL DEFINITION

As the markets consist of a diversity of components, different relaxation times occur because of the underlying relativistic effects between different components. It follows that the weak form of efficiency coupled with the rational expectation model cannot be attained. Because of the presence of different time components with heterogeneous expectations, current market prices cannot reflect all available information. The price discovery mechanism follows rather a dynamic “error correction model” where the successive reactions to an event unfold in the price. Why, then, did this not show up more clearly in previous scientific investigations? Some of the several reasons include the following:

- High-frequency data are a prerequisite for the empirical investigation of relativistic phenomena.
- Extensive computing power is needed to show the predictability in financial markets. Access to reasonably priced computing power has become available only recently.
- It is in the past few decades that an increasing awareness for dynamic and nonlinear processes has been gained. Such an awareness is crucial for the study of relativistic effects.

The presumption of conventional economics that forecasting is impossible per definition has had a powerful impact on the research on market efficiency. Economists have focused on structural studies that were hamstrung by a lack of high-frequency data and theoretical shortcomings. Little academic research has been invested in actually trying to predict shorter-term price movements and build successful trading models.

## 12.6 AN IMPROVED DEFINITION OF “EFFICIENT MARKETS”

Although the current definition of efficient markets has shortcomings, we do not think that this concept should be abandoned; rather, it should be adapted to the new findings. It is important to find a good measure of how well a market operates.

From a dynamic perspective, the notion of reduced friction should be central to the notion of efficiency. We consider an efficient market to be a market where all market information must be available to the decision makers and there must be participants with different time scales and heterogeneous expectations trading with each other to ensure a minimum of friction in the transaction costs.

A quantitative measure of efficiency might be derived from the bid-ask spreads (those between real bid and ask prices being more appropriate for such a measure

CONFIDENTIAL

than the nominal spreads quoted in information systems). Spreads are not only a measure of "friction," they also contain a risk component. The volatility or, more precisely, the probability of extreme returns within short time intervals should be considered together with the spread in the quantitative measure of market efficiency to be proposed. We are sure that in the years to come this definition will prevail and we shall find precise measures of efficiency as it is the case in thermodynamics and engineering.

FOR THE COURT

## BIBLIOGRAPHY

---

- Admati, A. R., and Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability, *Review of Financial Studies*, 1, 3–40.
- Ahn, D.-H., Boudoukh, J., Richardson, M., and Whitelaw, R. F. (2000). Partial adjustment or stale prices? Implications from stock index and futures return autocorrelations, *Stern Business School Working Paper*.
- Allais, M. (1974). The psychological rate of interest, *Journal of Money, Credit and Banking*, 3, 285–331.
- Allen, F., and Karjalainen, R. (1999). Using genetic algorithms to find technical trading rules, *Journal of Financial Economics*, 51, 245–271.
- Almeida, A., Goodhart, C. A. E., and Payne, R. G. (1998). The effects of macroeconomic news on high frequency exchange rate behavior, *Journal of Financial and Quantitative Analysis*, 33, 383–408.
- Andersen, T. G. (1996). Return volatility and trading volume: An information flow interpretation of stochastic volatility, *Journal of Finance*, 51, 169–204.
- Andersen, T. G., and Bollerslev, T. (1997a). Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns, *Journal of Finance*, 52, 975–1005.
- Andersen, T. G., and Bollerslev, T. (1997b). Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance*, 4, 115–158.
- Andersen, T. G., and Bollerslev, T. (1998a). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review*, 39, 885–905.

- Andersen, T. G., and Bollerslev, T. (1998b). Deutsche Mark-Dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance*, 53, 219–265.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2000). Exchange rate returns standardized by realised volatility are (nearly) Gaussian, *NBER Working Paper No: 7488*.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001). The distribution of realized exchange rate volatility, *Journal of the American Statistical Association*, forthcoming.
- Baestaens, D. J. E., and Van den Bergh, W. M. (1995). The marginal contribution of news to the DEM/USD swap rate, *Neural Network World*, 5, 371–381.
- Baillie, R. T., and Bollerslev, T. (1989). The message in daily exchange rates: A conditional-variance tale, *Journal of Business and Economic Statistics*, 7, 297–305.
- Baillie, R. T., and Bollerslev, T. (1990). Intraday and intermarket volatility in foreign exchange rates, *Review of Economic Studies*, 58, 565–585.
- Baillie, R. T., and McMahon, P. C. (1989) *The Foreign Exchange Market*, Cambridge University Press, Cambridge.
- Baillie, R. T., Bollerslev, T., and Mikkelsen, H.-O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics*, 74, 3–30.
- Balocchi, G. (1996). Interbank money market rate statistics: Results for the forecaster optimization, Internal document GBA.1996-03-01, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Balocchi, G., and Hopman, C. (1997), Bond futures: From statistical properties to forecasting, Internal document GBA.1997-10-28, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Balocchi, G., Dacorogna, M. M., Hopman, C. M., Müller, U. A., and Olsen, R. B. (1999a). The intraday multivariate structure of the Eurofutures markets, *Journal of Empirical Finance*, 6, 479–513.
- Balocchi, G., Dacorogna, M. M., Gençay, R., and Piccinato, B. (1999b). Intraday statistical properties of Eurofutures, *Derivatives Quarterly*, 6, 28–44.
- Balocchi, G., Gençay, R., Piccinato, B., and Dacorogna, M. M. (2001). Time-to-expiry seasonalities in Eurofutures, *Studies in Nonlinear Dynamics and Econometrics*, forthcoming.
- Bank for International Settlements (1990). Survey of foreign exchange market activity, *Document from the Monetary and Economic Department*, Basle, February 1990.
- Bank for International Settlements (1993). Central bank survey of foreign exchange market activity, *Document from the Monetary and Economic Department*, Basle, February 1993.

- Bank for International Settlements (1995). Central bank survey of foreign exchange market activity, *Document from the Monetary and Economic Department*, Basle, February 1995.
- Bank for International Settlements (1999). Central bank survey of foreign exchange and derivatives market activity, *Document from the Monetary and Economic Department*, Basle, May 1999.
- Banzhaf, W., Nordin, P., Keller, R. E., and Francone, F. D. (1998). *Genetic Programming: An Introduction*, Morgan Kaufmann, San Francisco, CA.
- Barndorff-Nielsen, O. E. (1998). Processes of normal inverse Gaussian type, *Finance and Stochastics*, 2, 41–68.
- Barndorff-Nielsen, O. E., and Prause, K. (1999). Apparent scaling, *Finance and Stochastics*, forthcoming.
- Bekaert, G., and Hodrick, R. J. (1992). Characterizing predictable components in excess returns on equity and foreign exchange markets, *Journal of Finance*, 47, 467–509.
- Beltratti, A., and Morana, C. (1998). Estimating variance in the foreign exchange market with high frequency data, *Bocconi University*, manuscript.
- Beltratti, A., and Morana, C. (1999). Computing value at risk with high frequency data, In *Journal of Empirical Finance*, 6, 431–455.
- Berndt, E., Hall, B., Hall, R., and Hausman, J. (1974). Estimation and inference in nonlinear structural models, *Annals of Economic and Social Measurement*, 3, 653–665.
- Bhattacharyya, S., Pictet, O. V., and Zumbach, G. (1998). Representational semantics for genetic programming based learning in high-frequency financial data, In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, Madison, Wisconsin, M. Kaufmann, Ed., 11–16.
- Bjorn, V. (1994). Optimal multiresolution decomposition of financial time-series, In *Proceedings of the Conference on Neural Networks in the Capital Markets*, Pasadena, November 1994.
- Black, F., and Scholes, M. (1973). The pricing of option and corporate liabilities, *Journal of Political Economy*, 81, 637–659.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics*, 31, 307–327.
- Bollerslev, T., and Domowitz, I. (1993). Trading patterns and prices in the interbank foreign exchange market, *Journal of Finance*, 48, 1421–1443.
- Bollerslev, T., and Melvin, M. (1994). Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis, *Journal of International Economics*, 36, 355–372.
- Bollerslev, T., Chou, R. Y., and Kroner, K. F. (1992). ARCH modeling in finance, *Journal of Econometrics*, 52, 5–59.

2020 RELEASE UNDER E.O. 14176



- Boothe, P., and Glassman, D. (1987). The statistical distribution of exchange rates, empirical evidence and economic implications, *Journal of International Economics*, 22, 297–319.
- Bouchaud, J.-P., Potters, M. (2000). Theory of financial risks, *From Statistical Physics to Risk Management*, Cambridge University Press, Cambridge.
- Bouchaud, J.-P., Potters, M., and Meyer, M. (2000). Apparent multifractality in financial time series, *European Physical Journal B*, 13, 595–599.
- Boyer, B. H., Gibson, M. S., and Loretan, M. (1997). Pitfalls in tests for changes in correlations, *Federal Reserve Board International Finance Discussion Paper No: 597*.
- Breymann, W. (2000). Dynamic theta time: Algorithm, configuration, tests, Internal document WAB.2000-07-31, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Breymann, W., Ghashghaie, S., and Talkner, P. (2000). A stochastic cascade model for FX dynamics, *International Journal of Theoretical and Applied Finance*, 3, 357-360.
- Brock, W., Lakonishok, J., and LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns, *Journal of Finance*, 47, 1731–1764.
- Brock, W. A. (1993). Pathways to randomness in the economy: Emergent nonlinearity and chaos in economics and finance, *Estudios Economicos*, 8, 3–55.
- Brock, W. A. (1999). Scaling in economics: A reader's guide, *Industrial and Corporate Change*, 8, 409–446.
- Brock, W. A., and Hommes, C. (1997). A rational route to randomness, *Econometrica*, 65, 1059–1095.
- Brock, W. A., and Kleidon, A. (1992). Periodic market closure and trading volume: A model of intraday bids and asks, *Journal of Economic Dynamics and Control*, 16, 451–489.
- Brock, W. A., and LeBaron, B. (1996). A dynamical structural model for stock return volatility and trading volume, *Review of Economics and Statistics*, 78, 94–110.
- Burghardt, G., and Hoskins, B. (1995). The convexity bias in Eurodollar futures: Part I, in *The Handbook on Derivative Instruments*, Atsuo Konishi and Ravi E. Dattareya, Eds., Ch. 4, 81–120.
- Calderon-Rossel, J. R., and Ben-Horim, M. (1982). The behavior of the foreign exchange rates, empirical evidence and economic implications, *Journal of International Business Studies*, 13, 99–111.
- Calvet, L. E., Fisher, A. J., and Mandelbrot, B. B. (1997). Large deviations and the distribution of price changes, *Cowles Foundation Working Paper No: 1165*.
- Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*, Princeton University Press, Princeton, New Jersey.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112  
113  
114  
115  
116  
117  
118  
119  
120  
121  
122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319  
320  
321  
322  
323  
324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863  
864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917  
918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971  
972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000

- Chopard, B., Oussaidène, M., Pictet, O., Schirru, R., and Tomassini, M. (1995). Evolutionary algorithms for multimodal optimization in financial applications, In *Proceedings of the SPP-IF Seminar, Zürich*, 139–142.
- Chopard, B., Pictet, O., and Tomassini, M. (2000). Parallel and distributed evolutionary computation for financial applications, *Parallel Algorithms and Applications Journal*, 15, 15–36.
- Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices, *Econometrica*, 41, 135–155.
- Cohen, K., Hawawimi, G., Maier, S., Schwartz, R., and Whitcomb, D. (1983). Friction in the trading process and the estimation of systematic risk, *Journal of Financial Economics*, 12, 263–278.
- Corsi, F., Zumbach, G. O., Müller, U. A., and Dacorogna, M. M. (2001). Consistent high-precision volatility from high-frequency data, *Economic Notes*, forthcoming.
- Cox, J. C., and Rubinstein, M. (1985). *Options Markets*, Prentice-Hall, Englewood Cliffs, NJ.
- 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”, July 27–July 31, 1992; MMD.1992-05-12, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Dacorogna, M. M., Gauvreau, C. L., Müller, U. A., Olsen, R. B., and Pictet, O. V. (1996). Changing time scale for short-term forecasting in financial markets, *Journal of Forecasting*, 15, 203–227.
- 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, 413–438.
- Dacorogna, M. M., Müller, U. A., Jost, C., Pictet, O. V., Olsen, R. B., and Ward, J. R. (1995). Heterogeneous real-time trading strategies in the foreign exchange market, *European Journal of Finance*, 1, 243–253.
- Dacorogna, M. M., Müller, U. A., Embrechts, P., and Samorodnitsky, G. (1996). How heavy are the tails of a stationary HARCH(k) process?, published in “Stochastic Processes and Related Topics” book in memory of Stamatis Cambanis edited by Yannis Karatzas Balram S. Rajput and Murad Taqqu published by Birkhäuser, Boston.
- Dacorogna, M. M., Embrechts, P., Müller, U. A., and Samorodnitsky, G. (1998a). How heavy are the tails of a stationary HARCH(k) process? In *Stochastic Processes and Related Topics*, Y. Rajput and M. Taqqu, Eds., pp. 1–31, Birkhäuser, Boston.
- Dacorogna, M. M., Müller, U. A., Olsen, R. B., and Pictet, O. V. (1998b). Modelling short-term volatility with GARCH and HARCH models, In *Nonlinear Modelling of High Frequency Financial Time Series*, C. Dunis and B. Zhou, Eds., pp. 161–176, John Wiley & Sons, Chichester.

- Dacorogna, M. M., Pictet, O. V., Müller, U. A., and de Vries, C. G. (2001a). Extremal forex returns in extremely large data sets, *Extremes*, forthcoming.
- Dacorogna, M. M., Gençay, R., Müller, U. A., and Pictet, O. V. (2001b). Effective return, risk aversion and drawdowns, *Physica A*, 289, 66-88.
- Danielsson, J., de Haan, L., Peng, L., and de Vries, C. G. (1997). Using a bootstrap method to choose the sample fraction in tail index estimation, *Journal of Multivariate Analysis*, forthcoming.
- Davé, R. D. (1993). Statistical correlation of data frequency price change and spread results, Internal document RDD.1993-04-26, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Davé, R. D., and Stahl, G. (1998). On the accuracy of var estimates based on the variance-covariance approach, In *Proceedings of the 6th Karlsruher Ökonometrie-Workshop*, March 1997, Springer Verlag.
- Davidson, R., and MacKinnon, J. G. (1993). *Estimation and Inference in Econometrics*, Oxford University Press, Oxford, England.
- Davis, R., Mikosch, T., and Basrak, B. (1999). Sample ACF of multivariate stochastic recurrence equations with applications to GARCH, *Department of Mathematics and Computing Science of the University of Groningen Working Paper*.
- de Haan, L. (1990). Fighting the ARCH-enemy with mathematics, *Statistica Neerlandica*, 44, 45-68.
- de Haan, L., Jansen, D. W., Koedijk, K. G., and de Vries, C. G. (1994). Safety first portfolio selection, extreme value theory and long run asset risks, In *Extreme value Theory and Applications*, J. Galambos, J. Lechner, and E. Simiu, Eds., pp. 471-488, Kluwer, Dordrecht.
- de Jong, F., and Nijman, T. (1997). High frequency analysis of lead-lag relationships between financial markets, *Journal of Empirical Finance*, 4, 259-277.
- de Vries, C. G. (1992). Stylized facts of nominal exchange rate returns, in *Handbook of International Macroeconomics*, F. van der Ploeg, Ed., pp. 1-57.
- Deb, K., and Goldberg, D. E. (1989). An investigation of niche and species formation in genetic function optimization, In *Proceeding of the Third International Conference on Genetic Algorithms*, pp. 42-50.
- Demos, A. A., and Goodhart, C. A. E. (1992). The interaction between the frequency of market quotations, spread, and volatility in the foreign exchange market, *Applied Economics*, 28, 28, 377-386.
- Diebold, F. X. (1988). Empirical modeling of exchange rate dynamics, *Lecture Notes in Economics and Mathematical Systems*, Vol. 3030, Springer Verlag, Berlin.
- 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.

- Drost, F., and Nijman, T. (1993). Temporal aggregation of GARCH processes, *Econometrica*, 61, 909–927.
- Dunis, C., and Feeny, M. (1989). *Exchange Rate Forecasting*, Woodhead-Faulkner, Cambridge.
- Dunis, C., and Keller, A. (1993). Implied versus historical volatility: An empirical test of the efficiency of the currency options market using non-overlapping data, *Conference on Financial Markets Dynamics and Forecasting*, Paris, 2–4 Sep 1993.
- Eben, K. (1994). Arbitrage alerts and changing interrelations between FX rates, *Institute of Computer Science, Academy of Sciences of the Czech Republic*, manuscript.
- Eberlein, E., Keller, U., and Prause, K. (1998). New insights into smile – mispricing and value at risk: The hyperbolic model, *Journal of Business*, 71, 371–405.
- Ederington, L. L., and Lee, J. H. (1993). How markets process information: News releases and volatility, *Journal of Finance*, 48, 1161–1191.
- Ederington, L. L., and Lee, J. H. (1995). The short-run dynamics of the price adjustment to new information, *Journal of Financial and Quantitative Analysis*, 30, 117–134.
- Edison, H. (1993). The effectiveness of central bank intervention: A survey of the post-1982 literature, *Essays in International Finance*, Princeton University.
- Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events, Applications of Mathematics Stochastic Modelling and Applied Probability*, Vol. 33, Springer, Berlin.
- Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation, *Econometrica*, 50, 987–1008.
- Engle, R. F., and Russell, J. R. (1997). Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model, *Journal of Empirical Finance*, 4, 187–212.
- Engle, R. F., and Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data, *Econometrica*, 66, 1127–1162.
- 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.
- Epps, T. (1979). Co-movements in stock prices in the very short-run, *Journal of the American Statistical Association*, 74, 291–298.
- Evetts, M., and Fernandez, T. (1998). Numeric mutation improves the discovery of numeric constants in genetic programming, In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, M. Kaufmann, Ed., pp. 66-71, Madison, Wisconsin.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work, *Journal of Finance*, 25, 383–417.

- Fama, E. F. (1991). Efficient capital markets: II, *Journal of Finance*, 46, 1575–1617.
- Farmer, J. D. (1998). Market force, ecology, and evolution, *Santa Fe Institute Working Paper*.
- Farmer, J. D., and Lo, A. W. (1999). Frontiers of finance: Evolution and efficient markets, *Proceeding of the National Academy of Science, U.S.A.*, 96, 9991–9992.
- Feinstone, L. J. (1987). Minute by minute: Efficiency, normality, and randomness in intradaily asset prices, *Journal of Applied Econometrics*, 2, 193–214.
- Feller, W. (1971). *An Introduction to Probability Theory and Its Applications*, Wiley Series in Probability and Mathematical Statistics, Vol. 2, 2nd ed., John Wiley & Sons, New York.
- Fisher, A. J., Calvet, L. E., and Mandelbrot, B. B. (1997). Multifractality of Deutschemark U.S. Dollar exchange rates, *Cowles Foundation Working Paper No: 1166*.
- Flood, M. D. (1991). Microstructure theory and the foreign exchange market, *Federal Reserve Bank of St. Louis Review*, 73, 52–70.
- Flood, M. D. (1994). Market structure and inefficiency in the foreign exchange market, *Journal of International Money and Finance*, 13, 131–158.
- Frances, P. H., and van Griensven, K. (1998). Forecasting exchange rates using neural networks for technical trading rules, *Studies in Nonlinear Dynamics and Econometrics*, 4, 109–114.
- Franke, G., and Hess, D. (1997). Information diffusion in electronic and floor trading, *University of Konstanz Working Paper*.
- Franke, G., and Hess, D. (1998). The impact of scheduled news announcements on T-bond and Bund futures trading, *University of Konstanz Working Paper*.
- Frankel, J. A. and Froot, K. A. (1990). Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review*, 80, 181–185.
- Friedmann, D. and Vandersteel, S. (1982). Short-run fluctuations in foreign exchange rates, evidence from the data 1973–79, *Journal of International Economics*, 13, 171–186.
- Frisch, U. (1995). *Turbulence*, Cambridge University Press, Cambridge, U.K.
- Froot, K. A. (1993). Currency hedging over long horizons, *National Bureau of Economic Research Working Paper No: 4355*.
- Froot, K. A., Kim, M., and Rogoff, K. (1995). The law of one price over 700 years, *National Bureau of Economic Research Working Paper No: 5132*.
- Fung, H.-G., and Leung, W. K. (1993). The pricing relationship of Eurodollar futures and Eurodollar deposit rates, *Journal of Futures Markets*, 13, 115–126.
- Gallant, A., Rossi, P., and Tauchen, G. (1993). Nonlinear dynamic structures, *Econometrica*, 61, 871–907.
- Garbade, K. D., and Silber, W. L. (1985). Price movements and price discovery in futures and cash markets, *Review of Economics and Statistics*, 65, 289–297.

2025 RELEASE UNDER E.O. 14176

- Gençay, R. (1998a). Optimization of technical trading strategies with neural network models and evidence of profitability in security markets, *Economics Letters*, 59, 249–254.
- Gençay, R. (1998b). The predictability of security returns with simple technical trading rules, *Journal of Empirical Finance*, 5, 317–345.
- Gençay, R. (1999). Linear, nonlinear and essential foreign exchange prediction with simple technical rules, *Journal of International Economics*, 47, 91–107.
- Gençay, R., and Stengos, T. (1998). Moving average rules, volume and the predictability of security returns with feedforward networks, *Journal of Forecasting*, 17, 401–414.
- Gençay, R., Selçuk, F., and Whitcher, B. (2001a). Differentiating intraday seasonalities through wavelet multiscaling, *Physica A*, 289, 543–556.
- Gençay, R., Selçuk, F., and Whitcher, B. (2001b). *An Introduction to Wavelets and Other Filtering Methods in Economics and Finance*, Academic Press, forthcoming, San Diego, CA.
- Gençay, R., Dacorogna, M. M., Olsen, R., and Pictet, O. V. (2001c). Real-time foreign exchange trading models and market behavior, *Journal of Economic Dynamics and Control*, forthcoming.
- Gençay, R., Selçuk, F., and Whitcher, B. (2001d). Scaling properties of foreign exchange volatility, *Physica A*, 289, 89–106.
- Gençay, R., Ballocci, G., Dacorogna, M. M., Olsen, R., and Pictet, O. V. (2002). Real-time trading models and the statistical properties of foreign exchange rates, *International Economic Review*, forthcoming.
- Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P., and Dodge, Y. (1996). Turbulent cascades in foreign exchange markets, *Nature*, 381, 767–770.
- Ghose, D., and Kroner, K. F. (1995). The relationship between GARCH and symmetric stable processes: Finding the source of fat tails in financial data, *Journal of Empirical Finance*, 2, 225–251.
- Ghysels, E. and Jasiak, J. (1995). Stochastic volatility and time deformation: An application of trading volume and leverage effects, In *Proceedings of the HFDF-I Conference*, Zurich, Switzerland, March 29-31, 1995, 1, 1–14.
- Ghysels, E., Harvey, A., and Renault, E. (1996). Stochastic volatility, in *Handbook of Statistics 14, Statistical Methods in Finance*, G. S. Mandala and C. S. Rao eds., North holland, Amsterdam.
- Gibson, M. S., and Boyer, B. H. (1997). Evaluating forecasts of correlation using option pricing, *Federal Reserve Board International Finance Discussion Paper Series No: 1997-600*.
- Gielens, G., Straetmans, S., and de Vries, C. G. (1996). Fat tail distributions and local thin tail alternatives, *Communications in Statistics, Theory and Methods*, 25, 705–710.
- Glassman, D. (1987). Exchange rate risk and transactions costs: Evidence from bid-ask spreads, *Journal of International Money and Finance*, 6, 479–490.

2025 RELEASE UNDER E.O. 14176

- Glosten, L. R. (1987). Components of the bid-ask spread and the statistical properties of transaction prices, *Journal of Finance*, 42, 1293–1307.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization & Machine Learning*, Addison-Wesley, Reading, MA.
- Goldberg, D. E., and Richardson, J. (1987). Genetic algorithms with sharing for multimodal function optimization, In *Proceeding of the Second International Conference on Genetic Algorithms*, pp. 41–49.
- Goldie, C. M., and Smith, R. L. (1987). Slow variation with remainder: Theory and applications, *Quarterly Journal of Mathematics, Oxford 2nd series*, 38, 45–71.
- Goodhart, C. A. E. (1989). News and the foreign exchange market, In *Proceedings of the Manchester Statistical Society*, pp. 1–79.
- Goodhart, C. A. E., and Curcio, R. (1991). The clustering of bid/ask prices and the spread in the foreign exchange market, *LSE Financial Market Group Discussion Paper No: 110*.
- Goodhart, C. A. E., and Demos, A. (1990). Reuters screen images of the foreign exchange market: The Deutschmark/Dollar spot rate, *Journal of International Securities Markets*, 4, 333–348.
- 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 Figliuoli, L. (1992). The geographical location of the foreign exchange market: A test of an “islands” hypothesis, *Journal of International and Comparative Economics*, 1, 13–27.
- Goodhart, C. A. E. and Hesse, T. (1993). Central bank FX intervention assessed in continuous time, *Journal of International Money and Finance*, 12, 368–389.
- Goodhart, C. A. E., and Payne, R. G. (1996). Microstructural dynamics in a foreign exchange electronic broking system., *Journal of International Money and Finance*, 15, 829–852.
- Goodhart, C. A. E., Hall, S. G., Henry, S. G. B., and Pesaran, B. (1993). News effects in a high frequency model of the Sterling-Dollar exchange rate, *Journal of Applied Econometrics*, 8, 1–13.
- Goodhart, C. A. E., Ito, T., and Payne, R. G. (1995). One day in June, 1993: A study of the working of Reuters 2000-2 electronic foreign exchange trading system, *NBER Technical Working Paper No: 179*.
- Gouree, J. K., and Hommes, C. H. (2000). Heterogeneous beliefs and the nonlinear cobweb model, *Journal of Economic Dynamics and Control*, 24, 761–798.
- Granger, C., and Ding, Z. (1995). Some properties of absolute return: An alternative measure of risk, *Annales d'Economie et de Statistique*, 40, 67–91.
- Granger, C. W. J., and Newbold, P. (1977). *Forecasting Economic Time Series*, Academic Press, London.
- Groenendijk, P. A., Lucas, A., and de Vries, C. G. (1996). Stochastic processes, nonnormal innovations, and the use of scaling ratios, In *Proceedings of the third*



- International Conference on Forecasting Financial Markets*, London March 27-29, 1996, 1, 1–38.
- Guillaume, D. M. (1994). On the trend-following behavior of intradaily foreign exchange market and its relationship with the volatility, *Center for Economic Studies, Catholic University of Leuven*, manuscript.
- Guillaume, D. M., Dacorogna, M. M., and Pictet, O. V. (1994). On the intradaily performance of GARCH processes, In *Proceedings of the First International Conference on High Frequency Data in Finance, HFDF-I*, Zürich, Switzerland.
- Guillaume, D. M., Dacorogna, M. M., Davé, R. D., Müller, U. A., Olsen, R. B., and Pictet, O. V. (1997). From the bird's eye to the microscope: A survey of new stylized facts of the intradaily foreign exchange markets, *Finance and Stochastics*, 1, 95–129.
- Gwilym, O. and Sutcliffe, C. (1999). *High-Frequency Financial Market Data*, Risk Books, London.
- Hall, P. (1982). On some simple estimates of an exponent of regular variation, *Journal of the Royal Statistical Society, Series B*, 44, 37–42.
- Hall, P. (1990). Using the bootstrap to estimate mean square error and select smoothing parameter in nonparametric problem, *Journal of Multivariate Analysis*, 32, 177–203.
- Hamilton, J. D. (1994). *Time Series Analysis*, Princeton University Press, Princeton, NJ.
- Hansen, L. P., and Hodrick, R. J. (1980). Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy*, 88, 829–853.
- Harris, L. (1987). Transaction data tests of the mixture of distributions hypothesis, *Journal of Financial and Quantitative Analysis*, 22, 127–141.
- Hartmann, P. (1998). Do Reuters spreads reflect currencies' differences in global trading activity, *Journal of International Money and Finance*, 17, 757–784.
- Hasbrouck, J. (1998). Security bid/ask dynamics with discreteness and clustering: Simple strategies for modeling and estimation, In *Proceedings of the HFDF-II Conference*, Zurich, 2, pp. 1–24.
- Hasbrouck, J. (1999). Trading fast and slow: Security market events in realtime, *Stern School of Business Working Paper*.
- Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution, *Annals of Statistics*, 3, 1163–1173.
- Hols, M. C., and De Vries, C. G. (1991). The limiting distribution of extremal exchange rate returns, *Journal of Applied Econometrics*, 6, 287–302.
- Hommel, C. (2000) Financial markets as nonlinear adaptive evolutionary systems, *Quantitative Finance*, 1, 149–167.
- Hosking, J. (1996). Asymptotic distribution of the sample mean, autocovariances and autocorrelations of long memory time series, *Journal of Econometrics*, 73, 261–284.



- Hsieh, D. A. (1988). The statistical properties of daily foreign exchange rates: 1974–1983, *Journal of International Economics*, 24, 129–145.
- Hull, J. C. (1993). *Options, Futures and other Derivative Securities*, Prentice-Hall International, Englewood Cliffs, NJ.
- Ibragimov, I., and Linnik, Y. V. (1971). *Independent and Stationary Sequences of Random Variables*, Wolters-Noordhoff, Groningen.
- International Monetary Fund, (1993). International capital markets. Part I: Exchange rate management and international capital flows, *World Economic and Financial Surveys*, Washington DC, April 1993.
- Ito, T., and Roley, V. V. (1987). News from the U.S. and Japan: Which moves the yen/dollar exchange rate? *Journal of Monetary Economics*, 19, 255–277.
- J. P. Morgan (1996). RiskMetrics — Technical document. J. P. Morgan and International Marketing – Reuters Ltd.
- Jones, C. M., Kaul, G., and Lipson, M. L. (1994). Transactions, volumes and volatility, *Review of Financial Studies*, 7, 631–651.
- Karpoff, J. M. (1987). The relation between price changes and trading volume: A survey, *Journal of Financial and Quantitative Analysis*, 22, 109–126.
- Keeney, R. L., and Raiffa, H. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley & Sons, New York.
- Kendall, M., Stuart, A., and Ord, J. K. (1987). *Advanced Theory of Statistics*, Vol. 1, 5th ed., Charles Griffin & Company Limited, London.
- Koedijk, K. G., Schaafgans, M. M. A., and De Vries, C. G. (1990). The tail index of exchange rate returns, *Journal of International Economics*, 29, 93–108.
- Koza, J. R. (1992). *Genetic Programming*, The MIT Press, Cambridge, MA.
- Krzanowski, W. J., and Marriott, F. H. C. (1994). *Multivariate Analysis, Distributions, Ordination and Inference*, Vol. 1, 1st ed., Kendall's Library Statistics, Edward Arnold, London.
- Krzanowski, W. J., and Marriott, F. H. C. (1995). *Multivariate Analysis, Classification, Covariance Structures and Repeated Measurements*, Vol. 2, 1st ed., Kendall's Library Statistics. Edward Arnold, London.
- Kurz, M. (1994). *Endogenous Economic Fluctuations, Studies in the Theory of Rational Beliefs*, Vol. 6, Studies in Economic Theory. Springer, Berlin.
- Leadbetter, M., Lindgren, G., and Rootzén, H. (1983). *Extremes and Related Properties of Random Sequences and Processes*, Springer Series in Statistics. Springer-Verlag, New York-Berlin.
- LeBaron, B. (1992a). Do moving average trading rule results imply nonlinearities in foreign exchange markets?, *University of Wisconsin-Madison, SSRI Working Paper*.
- LeBaron, B. (1992b). Forecast improvements using a volatility index, *Journal of Applied Econometrics*, 7, 137–149.
- LeBaron, B. (1992c). Some relations between volatility and serial correlations in stock market returns, *Journal of Business*, 65, 199–219.

- LeBaron, B. (1997). Technical trading rules and regime shifts in foreign exchange, *Advanced Trading Rules*, E. Acar and S. Satchell, Eds., Butterworth-Heinemann.
- LeBaron, B. (1999a). Technical trading rule profitability and foreign exchange intervention, *Journal of International Economics*, 49, 125–143.
- LeBaron, B. (1999b). Volatility persistence and apparent scaling laws in finance, *Brandeis University*, manuscript.
- LeBaron, B. (2000). Evolution and time horizons in an agent based stock market, *Brandeis University*, manuscript.
- Levich, R. M., and Thomas, L. R. (1993a). Internationally diversified bond portfolios: The merit of active currency risk management, *National Bureau of Economic Research Working Paper No.4340*.
- 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, 451–474.
- LIFFE (1995a). *Government Bond Futures*. The London International Financial Futures and Options Exchange, London, United Kingdom.
- LIFFE (1995b). *International Bond Market*. The London International Financial Futures and Options Exchange, London, United Kingdom.
- Lo, A., Mamaysky, H., and Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation, *Journal of Finance*, LV, 1705–1765.
- Lo, A. W., and MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies*, 1, 41–66.
- Lo, A. W., and MacKinlay, A. C. (1990a). An econometric analysis of nonsynchronous trading, *Journal of Econometrics*, 45, 181–211.
- Lo, A. W., and MacKinlay, A. C. (1990b). When are contrarian profits due to stock market overreaction? *Review of Financial Studies*, 3, 175–205.
- Longin, F. and Solnik, B. (1995). Is the correlation in international equity returns constant: 1960–1990? *Journal of International Money and Finance*, 14, 3–26.
- Loretan, M., and Phillips, P. C. B. (1994). Testing the covariance stationarity of heavy-tailed time series, *Journal of Empirical Finance*, 1, 211–248.
- Low, A., Muthuswamy, J., and Sarkar, S. (1996). Time variation in the correlation structure of exchange rates: High frequency analyses, In *Proceedings of the Third International Conference on Forecasting Financial Markets*, London, England, March 27–29, 1996, 1, 1–24.
- Lux, T., and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market? *Nature*, 397, 498–500.
- Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: A micro-simulation of interacting agents? *International Journal of Theoretical and Applied Finance*, 4, 1–20.

- Lyons, R. K. (1995). Test of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics*, 39, 321–351.
- Lyons, R. K. (1996a). Foreign exchange volume: Sound and fury signifying nothing? In *The Microstructure of Foreign Exchange Markets*, J. Frankel et al., Eds., University of Chicago Press, pp. 183–201.
- Lyons, R. K. (1996b). Optimal transparency in a dealership market with an application to foreign exchange, *Journal of Financial Intermediation*, 5, 225–254.
- Lyons, R. K. (1998). Profits and position control: A week of FX dealing, *Journal of International Money and Finance*, 17, 97–115.
- MacDonald, R., and Taylor, M. P. (1992). Exchange rate economics, *IMF Staff Papers*, 39, 1–57.
- Makridakis, S., Wheelwright, S. C., and McGee, V. E. (1983). *Forecasting Methods and Applications*, 2nd ed., John Wiley & Sons, New York.
- Mandelbrot, B. B. (1963). The variation of certain speculative prices, *Journal of Business*, 36, 394–419.
- Mandelbrot, B. B. (1971). When can price be arbitrated efficiently? A limit to the validity of the random walk and martingale model, *Review of Economics and Statistics*, 53, 225–236.
- Mandelbrot, B. B. (1972). Statistical methodology for nonperiodic cycles: From the covariance to R/S analysis, *Annals of Economic and Social Measurement*, 1, 259–290.
- Mandelbrot, B. B. (1983). *The Fractal Geometry of Nature*, W. H. Freeman and Company, New York.
- Mandelbrot, B. B. (1997). *Fractals and Scaling in Finance*, Springer Verlag.
- 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.
- Mandelbrot, B. B., Fisher, A. J., and Calvet, L. E. (1997). A multifractal model of asset returns, *Cowles Foundation Discussion Paper No: 1164*.
- Mansfield, P., Rachev, S. T., and Samorodnitsky, G. (1999). Long strange segments of a stochastic process and long range dependence, *Cornell University, School of Operations Research and Industrial Engineering, Technical Report No: 1252*.
- Mantegna, R. N., and Stanley, H. E. (1995). Scaling behavior in the dynamics of an economic index, *Nature*, 376, 46–49.
- Mantegna, R. N., and Stanley, H. E. (2000). *An Introduction to Econophysics — Correlations and Complexity in Finance*, Cambridge University Press, Cambridge.
- Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*, Wiley, Yale University Press (1970), New York.
- Markowitz, H. M. (1987). *Mean-Variance Analysis in Portfolio Choice and Capital Markets*, Basil Blackwell, Oxford, Cambridge, MA.

2025 RELEASE UNDER E.O. 14176

- Mason, D. M. (1982). Laws of large numbers for sums of extreme values, *Annals of Probability*, 10, 754–764.
- McCulloch, J. H. (1997). Measuring tail thickness in order to estimate the stable index  $\alpha$ : A critique, *Journal of Business and Economic Statistics*, 15(1), 74–81.
- McFarland, J. W., Petit, R. R., and Sung, S. K. (1982). The distribution of foreign exchange price changes: Trading day effects and risk measurement, *Journal of Finance*, 37, 693–715.
- McNeil, A. J. and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach, *Journal of Empirical Finance*, 7, 271–300.
- Meese, R. A., and Rogoff, J. (1983). Empirical exchange rate models of the seventies, do they fit out of sample? *Journal of International Economics*, 14, 3–24.
- Melvin, M. T., and Yin, X. (2000). Public information arrival, exchange rate volatility and quote frequency, *Economic Journal*, 110, 465–490.
- Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*, Springer Verlag, Heidelberg.
- Michalewicz, Z. (1994). *Genetic Algorithms + Data Structures = Evolution Programs*, Springer-Verlag, 2nd ed., Berlin.
- Mikosch, T. and Starica, C. (1999). Change of structure in financial time series, long range dependence and the GARCH model, *University of Groningen*, manuscript.
- Montana, D. J. (1995). Strongly typed genetic programming, *Evolutionary Computation*, 3, 199–230.
- Moody, J. and Wu, L. (1995). Statistical analysis and forecasting of high frequency foreign exchange rates, In *Proceedings of the HFDF-I Conference*, Zurich, Switzerland, March 29-31, 1995, 3, 29–31.
- Morgan Guaranty (1996). *RiskMetrics<sup>TM</sup> – Technical Document*, 4th ed., Morgan Guaranty Trust Company of New York, New York.
- Mosteller, F., and Tukey, J. W. (1977). *Data Analysis and Regression*, Addison-Wesley, Reading MA.
- Müller, U. A. (1991). Specially weighted moving averages with repeated application of the EMA operator, Internal document UAM.1991-10-14, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A. (1993). Statistics of variables observed over overlapping intervals, Internal document UAM.1993-06-18, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A. (1996). Generating a time series of fixed-period spot interest rates from interest rate futures, Internal document UAM.1996-04-19, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

- Müller, U. A. (1999). The O&A filter for data in finance, Internal document UAM.1999-04-27, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A. (2000). Volatility Computed by Time Series Operators at High Frequency, In *Statistics and Finance: An Interface*, W. S. Chan, W. K. Li and Howell Tong, Eds., Imperial College Press, London.
- Müller, U. A., and Sgier, R. G. (1992). Statistical analysis of intraday bid-ask spreads in the foreign exchange market, Internal document UAM.1992-04-10, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A., Dacorogna, M. M., Olsen, R. B., Pictet, O. V., Schwarz, M., and Morgeneegg, 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.
- Müller, U. A., Dacorogna, M. M., Davé, R. D., Pictet, O. V., Olsen, R. B., and Ward, J. R. (1993a). Fractals and intrinsic time — A challenge to econometricians, in “Erfolgreiche Zinsprognose”, Ed. by B. Lüthje, Verband öffentlicher Banken, Bonn 1994.
- Müller, U. A., Dacorogna, M. M., and Pictet, O. V. (1993b). A trading model performance measure with strong risk aversion against drawdowns, Internal document UAM.1993-06-03, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A., Dacorogna, M. M., and Pictet, O. V. (1995). The error of statistical volatility of intradaily quoted price changes observed over a time interval, Internal document UAM.1995-07-31, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V., and von Weizsäcker, J. E. (1997a). Volatilities of different time resolutions – analyzing the dynamics of market components, *Journal of Empirical Finance*, 4, 213–239.
- Müller, U. A., Dacorogna, M. M., and Lundin, M. C. (1997b). Currency overlay using trading models under exposure constraints, Internal document UAM.1997-07-06, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Müller, U. A., Dacorogna, M. M., and Pictet, O. V. (1998). Heavy tails in high-frequency financial data, In *A Practical Guide to Heavy Tails: Statistical Techniques for Analysing Heavy Tailed Distributions*, R. J. Adler, R. E. Feldman and M. S. Taqqu, Eds., pp. 55–77, Birkhäuser, Boston, MA.
- Murphy, J. J. (1986). *Technical Analysis of the Futures Markets*, New York Institute of Finance, Prentice-Hall, New York.
- Mussa, M. (1979). Empirical regularities in the behavior of exchange rates and theories of the foreign exchange market, *Carnegie-Rochester Series on Public Policy*, 11, 9–58.
- Neely, C., Weller, P., and Dittmar, R. (1997). Is technical analysis in the foreign exchange market profitable? A genetic programming approach, *Journal of Financial and Quantitative Analysis*, 32, 405–426.

2025 RELEASE UNDER E.O. 14176

- Neftci, S. N. (1991). Naive trading rules in financial markets and Wiener-Kolmogorov prediction theory: A study of technical analysis, *Journal of Business*, 64, 549–571.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach, *Econometrica*, 59, 347–370.
- Nelson, D. B., and Foster, D. P. (1994). Asymptotic filtering theory for univariate ARCH models, *Econometrica*, 62, 1–41.
- Oussaidène, M., Chopard, B., Pictet, O. V., and Tomassini, M. (1997). Parallel genetic programming and its application to trading model induction, *Parallel Computing*, 23, 1183–1198.
- Pecen, L., Ramešová, N., Pelikán, E., and Beran, H. (1995). Application of the GUHA method on financial data, *Neural Network World*, 5, 565–571.
- Peiers, B. (1997). Informed traders, intervention, and price leadership: A deeper view of the microstructure of the foreign exchange market, *Journal of Finance*, 52, 1589–1614.
- Pesaran, M. H., and Timmerman, A. (1992). A simple nonparametric test of predictive performance, *Journal of Business and Economic Statistics*, 10, 461–465.
- Peters, E. E. (1989). Fractal structure in the capital markets, *Financial Analysts Journal*, 1989, 32–37.
- Peters, E. E. (1991). *Chaos and Order in Capital Markets*, A Wiley Finance Edition, John Wiley & Sons, New York.
- Petersen, M. A. and Fialkowski, D. (1994). Posted versus effective spreads, good prices or bad quotes? *Journal of Financial Economics*, 35, 269–292.
- Piccinato, B., Balocchi, G., and Dacorogna, M. M. (1997). A closer look at the Eurofutures market: Intraday statistical analysis, Internal document BPB.1997-08-25, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Pictet, O. V., Dacorogna, M. M., Müller, U. A., Olsen, R. B., and Ward, J. R. (1992). Realtime trading models for foreign exchange rates, *Neural Network World*, 2, 713–744.
- Pictet, O. V., Dacorogna, M. M., Chopard, B., Oussaidène, M., Schirru, R., and Tomassini, M. (1995). Using genetic algorithms for robust optimization in financial applications, *Neural Network World*, 5, 573–587.
- Pictet, O. V., Dacorogna, M. M., and Müller, U. A. (1998). Hill, bootstrap and jackknife estimators for heavy tails, In *A practical guide to heavy tails: Statistical Techniques for Analysing Heavy Tailed Distributions*, R. J. Adler, R. E. Feldman and M. S. Taqqu, Eds., pp. 283–310, Birkhäuser, Boston, MA.
- Poterba, J. M., and Summers, L. H. (1988). Mean reversion in stock prices: Evidence and implications, *Journal of Financial Economics*, 22, 27–59.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1986). *Numerical Recipes. The Art of Scientific Computing*, Cambridge University Press, Cambridge.

2025 RELEASE UNDER E.O. 14176

- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C. The Art of Scientific Computing*, Cambridge University Press, Cambridge.
- Priestley, M. B. (1989). *Nonlinear and Nonstationary Time Series Analysis*, Academic Press, London.
- Rogalski, R. J., and Vinso, J. D. (1978). Empirical properties of foreign exchange rates, *Journal of International Business Studies*, 9, 69–79.
- Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance*, 39, 1127–1139.
- Roy, A. D. (1952). Safety first the holding of assets, *Econometrica*, 20, 431–449.
- Rydberg, T. H., and Shephard, N. (1998). Dynamics of trade-by-trade price movements: Decomposition and models, *Aarhus Center for Analytical Finance Working Paper Series No: 21*.
- Schinasi, G. J., and Swamy, P. A. V. B. (1989). The out-of-sample forecasting performance of exchange rate models when coefficients are allowed to change, *Journal of International Money and Finance*, 8, 375–390.
- Schnidrig, R. and Würtz, D. (1995). Investigation of the volatility and autocorrelation function of the USD/DEM exchange rate on operational time scales, In *Proceedings of the HFDF-I Conference*, Zurich, Switzerland, March 29-31, 1995, 3, 1–19.
- Sentana, E. (1991). Quadratic ARCH models: A potential reinterpretation of ARCH models, *LSE Financial Markets Group Discussion Paper No: 122*.
- Sharpe, W. F. (1966). Mutual fund performance, *Journal of Business*, 39, 119–138.
- Sharpe, W. F. (1994). The Sharpe Ratio, *Journal of Portfolio Management*, 21, 49–59.
- Shiller, R. J. (1989). *Market Volatility*, The MIT Press, Cambridge, MA.
- Shiller, R. J. (2000). *Irrational Exuberance*, Princeton University Press, Princeton.
- Stock, J. H. (1988). Estimating continuous-time processes subject to time deformation, *Journal of the American Statistical Association*, 83, 77–85.
- Subrahmanyam, A. (1991). Risk aversion, market liquidity, and price efficiency, *Review of Financial Studies*, 4, 417–441.
- Sullivan, R., Timmermann, A., and White, H. (1999). Data-snooping, technical trading rule performance and the bootstrap, *Journal of Finance*, 54, 1647–1692.
- Surajaras, P., and Sweeney, R. J. (1992). *Profit-Making Speculation in Foreign Exchange Markets, The Political Economy of Global Interdependence*, Westview Press, Boulder, CO.
- Suvanto, A. (1993). Foreign exchange dealing, Essays on the Microstructure of the Foreign Exchange Market, *ETLA, The Research Institute of the Finnish Economy Working Paper A19*.
- Svensson, L. E. (1992). An interpretation of recent research on exchange rate target zones, *Journal of Economic Perspectives*, 6, 119–144.



- Swamy, P. A. V. B., and Schinasi, G. J. (1989). Should fixed coefficients be re-estimated every period for extrapolation? *Journal of Forecasting*, 8, 1–17.
- Tauchen, G. E., and Pitts, M. (1983). The price variability-volume relationship on speculative markets, *Econometrica*, 51, 485–505.
- Taylor, M. P., and Allen, H. (1992). The use of technical analysis in the foreign exchange market, *Journal of International Money and Finance*, 11, 304–314.
- Taylor, S. J. (1986). *Modelling Financial Time Series*, John Wiley & Sons, Chichester.
- Taylor, S. J. (1994). Modeling stochastic volatility: A review and comparative study, *Mathematical Finance*, 4, 183–204.
- Taylor, S. J. and Xu, X. (1997). The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance*, 4, 317–340.
- Tucker, A. L. and Scott, E. (1987). A study of diffusion processes for foreign exchange rates, *Journal of International Money and Finance*, 6, 465–478.
- Walmsley, J. (1992). *The Foreign Exchange and Money Market Guide*, Wiley Finance Edition, John Wiley & Sons, New York.
- Wasserfallen, W. (1989). Flexible exchange rates: A closer look, *Journal of Monetary Economics*, 23, 511–521.
- Wasserfallen, W., and Zimmermann, H. (1985). The behavior of intradaily exchange rates, *Journal of Banking and Finance*, 9, 55–72.
- Westerfield, R. (1997). The distribution of common stock price changes: An application of transactions time and subordinated stochastic models, *Journal of Financial and Quantitative Analysis*, 12, 743–765.
- White, H. (1980). A heteroscedasticity-consistent covariance matrix and a direct test for heteroscedasticity, *Econometrica*, 48, 421–448.
- Yin, X., and Gernay, N. (1993). A fast genetic algorithm with sharing scheme using cluster analysis methods in multimodal function optimization, In *Proceedings of International Conference on Artificial Neural Nets and Genetic Algorithms*, Innsbruck, Austria, pp. 450–457.
- Zhou, B. (1993). Forecasting foreign exchange rates subject to de-volatilization, *MIT Sloan School Working Paper No: 3510*.
- Zimmermann, H. J. (1985). *Fuzzy Set Theory and Its Applications*, Kluwer-Nijhoff Publishing, Boston, Dordrecht, Lancaster.
- Zumbach, G. (1996). Time series and operators, Internal document GOZ.1996-05-23, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Zumbach, G. (1997). Measure of association, Internal document GOZ.1997-11-21, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.
- Zumbach, G. (2000). The pitfalls in fitting GARCH processes, In *Advances in Quantitative Asset Management*, C. Dunis, Ed., Kluwer, Amsterdam.





# INDEX

## A

activity  
 activity variable, *see* variable  
 market activity, 175, 177, 183  
 approach  
 macroeconomic, 5  
 microstructure, 2  
 time series, 2  
 arbitrage  
 formula, 29  
 opportunities, 28  
 riskfree, 22  
 triangular arbitrage, 118, 127  
 ARCH, *see* model

## B

Bank for International Settlements, 137  
 basis point, 23, 125, 171  
 BHHH algorithm, 223  
 bias, 44  
 bias of realized volatility, *see* volatility  
 market maker bias, 154  
 bid-ask bounce, 124  
 Black and Scholes, 43  
 Brownian motion, 49  
 build-up, 52, 87, 223, 263  
 error, 56  
 time interval, 56

business-time scale, *see* time scale

## C

call for margin, 12  
 capital management system, 296  
 cash interest rates, 21  
 cheapest-to-deliver, 29, 31  
 conditional heteroskedasticity, 204  
 conditional predictability, 215  
 convergence, 163, 320  
 model estimation, 223  
 correlation, 268  
 adjusted correlation measure, 272  
 breakdown, 270, 277  
 covolatility weighting, 268, 270  
 Epps effect, 269, 288  
 linear correlation coefficient, 272  
 memory, 293  
 variations, 278  
 counterparty default risk, 11  
 covariance, 268  
 cross-covariance, 124  
 matrix, 116, 346  
 realized, 50  
 credit ratings, 28  
 credit risk, 128  
 credit spread, 21  
 cross rate, 16, 19  
 cross-covariance, 124

cumulative distribution function, 173

## D

daily time series, 174

data

- ask, 12, 15
- bid, 12, 15
- bid-ask price, *see* price
- bid-ask spread, *see* spread
- daily, 14
- data cleaning, *see* filter
- data filter, *see* filter
- dependent quote, 95
- effective price, *see* price
- effective spread, *see* spread
- empirical, 132
- high-frequency, 1, 10, 32, 33
- historical, 16
- homogeneous, 35
- inhomogeneous, 35
- interest rate, 21
- intraday, 14
- irregularly spaced, 1
- low-frequency, 14
- quote, 1
- quoted spread, *see* spread
- quotes, 15
- regularly spaced, 1
- sparse, 3
- synthetic regular, 54
- tick, *see* tick, 38
- tick-by-tick, 1, 4, 6, 51
- time stamp, 15
- traded spread, *see* spread
- transaction price, 1, 15
- weekly, 14

data cleaning, 82

data error, 85

- decimal error, 85
- human error, 85
  - intentional error, 85
  - unintentional error, 85
- repeated ticks, 86
- scaling problem, 86
- system error, 85
- test, 85
  - early morning test, 86
  - monotonic series, 86
- tick copying, 86

data providers, 11

- Bloomberg, 11, 15
- Bridge, 11, 46, 186
- Knight Ridder, 15, 186
- Reuters, 11, 15, 46, 186

Telerate, 15, 128, 186

data-generating process, 122, 328, 338, *see* model

day-of-the-week effect, 169

daylight saving time, 164, 190

delivery risk, 170

deseasonalized returns, 197

directing process, 176

direction change indicator, 47

direction quality, 262

distribution, *see* probability distribution

- nonstable, 142

Dow Jones Industrials, 8

dynamic memory, 255

dynamic optimization, 246

dynamic overlay, *see* hedging

## E

economic forecast, 129

economic news announcement, 130

econophysicist, 9

effective news, 130

effective number of observations, 50

efficient frontier, 341

efficient markets, 349, 354

electronic order-matching system, 15

- Electronic Broking Services (EBS), 15

- Reuters Dealing 2000, 15

EMA-HARCH, *see* model

Epps effect, *see* correlation, 293

equity indices, 32

EUREX, 24

EURIBOR, 25

Euro, 24

Eurodeposits, 21

Eurofutures contracts, 121

Euroaira, *see* market, Eurofutures

European Monetary System (EMS), 127

Euroyen, *see* market, Eurofutures

expiry, 12, 29, 31

- quarterly expiry, 24

- time-to-expiry, 31

exponential attenuation, 293

exponential decline, 209

exponential memory, 282

exposure, 339

extreme events, 6

extreme risks, 144

extreme value theory, 138

## F

fat-tailedness, 132

FIGARCH, *see* model

- filter
- adaptive method, 98
  - after-jump algorithm, 105
  - artificial quote method, 118
  - ask quote, 110
  - bid quote, 110
  - bid-ask quote, 110
  - bid-ask spread, 110
  - build-up period, 87
  - credibility, 93, 114
  - data cleaning, 82
  - data filter, 82
  - decimal error, 113
  - domain error, 111
  - filter configuration, 113
  - filter parameters, 116
  - filtering algorithm, 89
    - credibility, 89
    - full-tick filtering window, 89
    - scalar filtering window, 89
    - univariate filter, 89
  - filtering hypothesis, 113
    - error hypothesis, 113
    - winning hypothesis, 113
  - forward premiums/discounts, 111
  - full-quote filtering window, 109
    - quote splitting, 110
  - high-quality data, 100
  - historical mode, 115
  - historical operation, 87
  - interest rate, 111
  - level filter, 88, 91
  - level quote, 110
  - multivariate filtering, 100, 116
    - filtering sparse data, 116
  - next point interpolation, 96
  - pair filtering, 88, 93, 98
  - price, 111
  - real-time mode, 115
  - real-time operation, 87
  - repeated quotes, 100
  - scalar filtering window, 103
    - filter test, 104
    - the normal update, 104
  - scalar quote, 110
  - scalar window
    - dismissing scalar quotes, 107
  - scaling factor, 113
  - second scalar window, 108
  - sensitivity analysis, 120
  - short-term interest-rate futures, 111
  - spread filter, 98
  - spread quoting, 98
  - strong filter, 116
  - time scale, 100
  - timing, 87
  - trust capital, 90, 104
  - univariate filtering, 113, 116
  - validity test, 110
    - weak (tolerant) filter, 116
  - finite variance, 132
  - first position, 30
  - forecasting, 248
    - forecast accuracy, 246, 262
    - forecast effectiveness, 264
    - forecast horizon, 264
    - forecast quality, 261
    - forecasting model, 249
    - multivariate forecasting, 249
    - real-time price forecasting system, 250
    - signal correlation, 263
      - forecasting signal, 263
      - real signal, 263
    - volatility forecast, 250
  - forecasting performance, 243
    - benchmark comparison, 245
    - direction quality, 245
    - realized potential, 245
  - foreign exchange (FX) market, *see* market
  - forward discount, 23
  - forward interest rate, 25
  - forward points, 23
  - forward premium, 23
  - forward rate, 22, 25
  - fractal behavior, 8, 209
  - fractional noise process, 207
  - FXFX page, 16
- G**
- GARCH, *see* model
  - Gaussian distribution, 135, 155
  - genetic algorithm, 223, 317
    - adaptive clustering, 319
    - multi-modal function, 318
    - sharing scheme, 319
  - genetic programming, 311
    - closure property, 312
    - function nodes, 312
    - syntactic restrictions, 313
    - terminal nodes, 312
    - tournament selection, 315
  - geographical components, 179
  - geometric mean, 170
  - goodness-of-fit, 285
  - granularity, 171
- H**
- HARCH, *see* model

- heat-wave component, 207
  - heat-wave effect, 209, 214, 352
  - hedge funds, 17
  - hedging, 31
    - currency overlay, 343
    - currency risk, 209, 340
    - dynamic hedging, 340, 345
    - instruments, 24
    - neutral point, 343
    - ratio, 343
  - heterogeneity, 44
  - heterogeneous market hypothesis, *see* hypothesis
  - heteroskedasticity, 35, 162
    - autoregressive conditional, 265
  - Hill estimator, 145
  - holiday
    - half-day, 190
  - holidays, *see* market
  - hyperbolic decline, 210
  - hypothesis
    - heat wave hypothesis, 179
    - heterogeneous market hypothesis, 209, 210, 224
    - island hypothesis, 179
    - meteor shower hypothesis, 179, 206
- I**
- IGARCH, *see* model
  - implied interest rate, 25
  - implied volatility, *see* volatility
  - index
    - AMEX Stock Index, 283
    - Down Jones Index, 283
  - indicator, 257
    - antisymmetric, 315
    - cycle, 310
    - overbought and oversold, 310
    - symmetric, 315
    - timing, 310
    - trend following, 310
    - volatile indicator, 259
  - information set, 82, 249
  - instability, 259
  - institutional constraints, 14, 128
  - institutional framework, 127
  - institutional investors, 17
  - interbank interest rates, 21
  - interbank money market rates, 121
  - interpolation, *see* method
  - intervention, 129
    - official, 129
  - intraday
    - analysis, 127
    - movements, 174
    - prices, 174
    - statistics, 163
    - volatility, 174
  - intraday
    - analysis, 177
    - statistics, 163
    - volatility, 174
  - intrinsic time, *see* time scale,  $\tau$ -scale
  - investment assets, 339
- J**
- J. P. Morgan, 6
- K**
- kernels, 52
  - Kronecker symbol, 260
  - kurtosis, 134, 205
    - operator, 71
- L**
- lagged correlation, 211
  - lead-lag, 20
  - lead-lag correlation, 211
  - leptokurticity, 173
  - LIBOR, 25
  - LIFFE, 24, 31
  - likelihood
    - likelihood-ratio test, 230
    - log-likelihood, 222
    - maximum likelihood, 223, 330
  - Ljung-Box, 334
  - long memory, 8, 198, 207
  - long-term regime, 259
- M**
- mapping function, 261
  - mark-to-market, 305
  - market
    - Asian stock, 52
    - bond, 28
    - centralized, 11
    - decentralized (OTC), 11
    - derivative, 11
    - equity, 32
    - Eurofutures, 169
      - Eurodollar, 169
      - Euroaira, 24, 169
      - Euromark, 169
      - Euroyen, 24

2025 RELEASE UNDER E.O. 14176

- Short Sterling, 169
- foreign exchange (FX), 11, 13, 13
  - FX forward, 15
  - FX forwards, 14
  - FX futures, 14
  - FX spot, 13, 15
- futures, 11, 12
  - bond futures, 28
  - commodity futures, 31
  - Eurofutures, 24
  - individual equity futures, 33
- geographical market, 214
- German bond market, 131
- heterogeneous market, 210
- holidays, 190
- homogeneous market, 210
- interest rate futures, 23, 25
- liquid, 1
- market microstructure, 5
- opening hours, 176
- option, 11, 13, 33
- over-the-counter (OTC), 2, 12, 19, 21, 24
- over-the-counter interest rate, 23
- participants, 10
  - anonymity, 10
- spot, 11, 12
- spot interest rate, 24
- U.S. treasury bond market, 131
- market activity, 175, 177, 183
- market efficiency, 14, 45, 156, 249, *see efficient markets*
- market expectation, 130
- market makers, 19
- market microstructure effect, 197, 201
- market risk, 158
- market-dependent persistence, 206
- Markov chain, 233
- maturity, 22
- maximum likelihood, *see likelihood*
- mean absolute error (MAE), 262
- mean square error (MSE), 262
- mean squared prediction errors (MSPE), 295
- measure
  - asymmetric effective returns, 307
  - reward-to-variability ratio, 305
  - risk-sensitive performance, 304
  - symmetric effective returns, 305
- method
  - distribution free measure, 262
  - interpolation, 37
    - linear, 54
    - previous-tick, 37
  - nonparametric method, 262
  - overlapping, 44
  - panel regression, 47
  - polynomial, 25
  - polynomial interpolation, 26
  - microstructure, 2, 5, 14
  - Middle European Time (MET), 326
  - middle price, 122
  - misspecification, 222
  - mixture of distributions, 131
  - model
    - ARCH, 221
    - capital asset pricing model (CAPM), 349
    - EMA-HARCH, 237, 254
    - FIGARCH, 221, 231
    - GARCH, 44, 146, 221, 222, 224, 328
      - AR-GARCH, 328
      - diffusion process, 224
      - estimation problems, 226
      - jump process, 224
      - temporal aggregation, 224
    - HARCH, 146, 231, 335
      - AR-HARCH, 328
      - HARCH components, 236
    - IGARCH, 251
    - in-sample optimization, 263
    - intraday market activity, 179
    - lagged adjustment model, 124
    - macroeconomic, 5
    - market activity, 175, 183
    - market maker bias, 154
    - market microstructure, 14
    - model initialization, 263
    - model structure, 256
    - moving average model, 295
    - multi-horizon, 263
    - multicascade model, 148
    - multifractal model, 148
    - multivariate volatility model, 250
    - nonparametric conditional mean models, 295
    - Normal Inverse Gaussian (NIG) Lévy process, 151
    - out-of-sample test, 263
    - purchasing power parity model, 249
    - QARCH, 232
    - random walk, 54, 147, 150, 152, 253, 331
    - risk premia, 249
    - structural model, 249
    - time series model, 249
    - trading, *see trading model*
    - trading model, 295
    - volatility, 6
- moment, 55, 132, 135, 137
  - finite second moment, 142
  - nonconverging fourth moment, 142
- momenta, 257
- Monte Carlo simulations, 145

multifractal model, 8  
 multifractality, 148  
 multiple assets, 278

**N**

near-singularity, 259  
 neutral point, *see* hedging  
 noise trader, 349  
 nonstable distribution, 142  
 notional deposit, 26

**O**

official intervention, 129  
 Olsen & Associates (O&A), 4, 11, 18  
 open position (mark-to-market), 305  
 opening hours, *see* trading hours  
 operator, 35  
   average, 55  
   backward shift, 77  
   causal, 54  
   comparison, 314  
   complex moving average, 75  
   conditional, 314  
   convolution, 51  
   crossover, 312  
   derivative, 55, 66  
   difference, 78  
   differential, 64  
   exponential moving average (EMA), 59  
     iterated, 59  
   homogeneous, 58  
   linear operator and kernels, 54  
   logic, 314  
   microscopic, 36, 76  
   moving average (MA), 61  
   moving correlation, 71  
   moving kurtosis, 71  
   moving norm, 63  
   moving skewness, 71  
   moving standard deviation, 63  
   moving variance, 63  
   mutation, 312  
   nonlinear, 58  
   regular time series, 77  
   tick frequency, 79  
   time translation, 77  
   time-translation invariant, 54  
   volatility, 68, 79  
   windowed Fourier transform, 74  
 optimization  
   robust optimization, 317  
 option pricing, 3  
 out-of-sample, 222

outright forward rate, 22  
 outright forward transactions, 23  
 overlap  
   method, 44  
   overlap-free, 50  
   overlapping returns, *see* return  
 overshooting, 27

**P**

performance index, 32  
 periodicity, 160, *see* seasonality  
 permanence hypothesis, 5  
 perpetuum mobile, 353  
 physical delivery, 31  
 portfolio  
   pricing, 277  
   trading model portfolio, 338  
 post ex-ante testing, 264  
 prefiltering technique, 249  
 price, 37, 38  
   bid-ask, 17  
   effective price, 40, 125  
   middle, 122  
   price formation, 3, 123  
   synthetic, 275  
   transaction, 1, 15  
 probability density function (pdf), 54  
 probability distribution, 54, 132  
 process, *see* model  
 psychological time, 176

**Q**

quadratic programming, 346  
 quote, *see* data

**R**

random walk, *see* model  
 rational expectations, 210  
 real market value, 210  
 realized volatility, *see* volatility  
 regime shifts, 8  
 relaxation time, 352  
 return, 37, 40  
   deseasonalized, 197  
   mean, 48  
   nonoverlapping, 50  
   overlapping, 47  
   synthetic, 275  
 Reuters, 1, 21  
 Reuters Instrument Code (RIC), 15  
 risk management, 3, 14, 52, 277

risk profiles, 14  
 risk-sensitive measures, 305  
 RiskMetrics, 251  
   methodology, 251  
   volatility, 253  
 robustness, 58  
 rolling over, 13  
 rollover scheme, 29  
 root mean squared error (RMSE), 186

**S**

safety margins, 171  
 scaling laws, *see* stylized facts  
   empirical, 177  
 scaling properties, 8  
   significance, 9  
 seasonality, 35, 44, 175  
   daily seasonality, 204  
   geographical seasonality, 181  
   ordinary seasonality, 214  
   seasonal heteroskedasticity, 175, 265  
   seasonal volatility, 174  
   weekly seasonality, 204  
 second position, 30  
 segmentation, 128  
 serial expiry contracts, 24  
 settlement rules, 12  
 Sharpe ratio, 305  
 short memory, 207  
 signal processing, 338  
 Singapore International Monetary Exchange,  
   24  
 skewness, 173  
   operator, 71  
   realized, 46  
 slippage, 32  
 speculative bubbles, 349  
 spot  
   interest rates, 21  
   market, *see* market  
   trading, 11  
   transaction, 23  
 spread  
   average, 46, 172  
   bid-ask, 17, 19, 23, 40, 45, 91, 124, 125,  
     154, 170, 354  
   credit spread, 21  
   effective, 126  
   log, 45  
   quoted, 171  
   relative, 45  
   traded, 171  
 spurious persistence, 194  
 stability, 143

standard limit theory, 153  
 stochastic process, *see* model  
 stock indices, 32  
 stock splits, 32  
 stop-loss deal, 301  
 stop-profit algorithm, 299  
 stringent filter, 302  
 structural change, 259  
 stylized facts, 2, 14, 121, 127  
   autocorrelation of return, 121  
   autocorrelation study, 161  
   bid-ask bounce, 124  
   bid-ask spread, 170  
   daily and weekly patterns, 122  
   deterministic volatility, 169  
   discreteness, 125  
   distribution of returns, 122  
     fat-tailed, 122  
   distributional issues, 121  
   distributional properties, 132  
     bounded distributions, 135  
     fat-tailed distributions, 135  
     thin-tailed distributions, 135  
   negative first-order autocorrelation of re-  
     turns, 123  
   scaling laws, 122, 147  
     apparent scaling, 152  
     fat tails, 142  
     interquartile range, 150  
     limitations, 158  
   scaling properties, 121  
   seasonal heteroskedasticity, 122, 163  
   seasonal volatility, 163, 167  
     U-shaped, 167  
   seasonality, 121  
 subordinated process, 176  
 swap, 23  
   FX swap rates, 23  
 symmetry, 132  
 synthetic price, *see* price  
 synthetic return, *see* return

**T**

tail  
   index, 6, 132, 135  
   statistics, 6  
 Taylor expansion, 62  
 technical analysis, 3, 14, 129  
 technological change, 8  
 temporal aggregation, 224  
 test  
   ex-ante test, 322  
   likelihood ratio, 7  
   Monte Carlo, 274



out-of-sample test, 249  
 ratio test, 192  
 tick, 1, 10  
   tick frequency, 37, 46  
     log, 46  
   tick time, 124  
 tick-by-tick, *see* data  
 time horizon, 209, 210, 249, 307, 309  
 time scale, 8, 176  
    $\tau$ -scale, 255  
    $\vartheta$ -scale, *see* time scale,  $\vartheta$ -time  
    $\vartheta$ -time, 174, 176, 188  
   business, 174, 188, 225  
   tick time, 124  
   variety of time scales, 8  
 time stamp, 10, 15  
 trading horizon (trader class), 198  
   day trader, 199  
   intraday trader, 199  
   long-term trader, 199  
   market maker, 199  
   short-term trader, 199  
 trading hours, 174, 299, 327  
 trading model, 295  
   current price, 302  
   current return, 300  
   gearing calculator, 301  
   market constraints, 299  
   performance calculator, 297  
   performance measures, 304  
      $R_{eff}$ , 307  
      $X_{eff}$ , 305  
     Sharpe ratio, 305  
   portfolios, 338  
   real-time trading models, 296, 310, 323  
   real-time trading strategies, 297  
   recommendation maker, 301  
   simulated trader, 297, 303  
     bookkeeper, 303  
     opportunity catcher, 303  
   stop-loss deal, 301  
   stop-loss detector, 302  
   stop-loss price, 302  
   stop-profit algorithm, 299  
   stop-profit control, 302  
   symmetry properties, 315  
   testing procedures, 317  
   trading hours, *see* trading hours  
   transaction costs, *see* transaction  
 transaction  
   clock, 176  
   costs, 170, 295, 325, 331, 333  
   price, 1, 15  
   volume, 46, 176  
 trust capital, *see* filter

**U**

uncertainty, 154  
 undershooting, 28

**V**

Value-at-Risk (VaR), 6, 250  
 variable  
   activity, 176  
   direction change indicator, 47  
   price, *see* price  
   realized covariance, *see* covariance  
   realized skewness, *see* skewness  
   return, *see* return  
   spread, *see* spread  
   tick frequency, *see* tick  
   volatility, *see* volatility  
 vehicle currency, 19, 172  
 volatility  
   annualized, 41  
   coarse volatility, 211  
   conditional heteroskedasticity, 198  
   daily volatility, 158  
   deterministic volatility, 170  
   expected volatility, 96  
   fine volatility, 211  
   historical, 41  
   implied, 43  
   model, 43  
   patterns, 176  
   realized volatility, 37, 41, 197, 248  
     bias, 154, 159, 198  
     bias correction, 202  
   volatility clustering, 122, 161, 198, 228  
   volatility ratio, 47, 192

**W**

wavelet, 193  
   multiscaling approach, 193  
   signal-to-noise ratio, 193  
   wavelet transform, 159  
   wavelet variance, 160  
 weekend effect, 172  
 White's variance-covariance matrix estimation,  
   224

**Y**

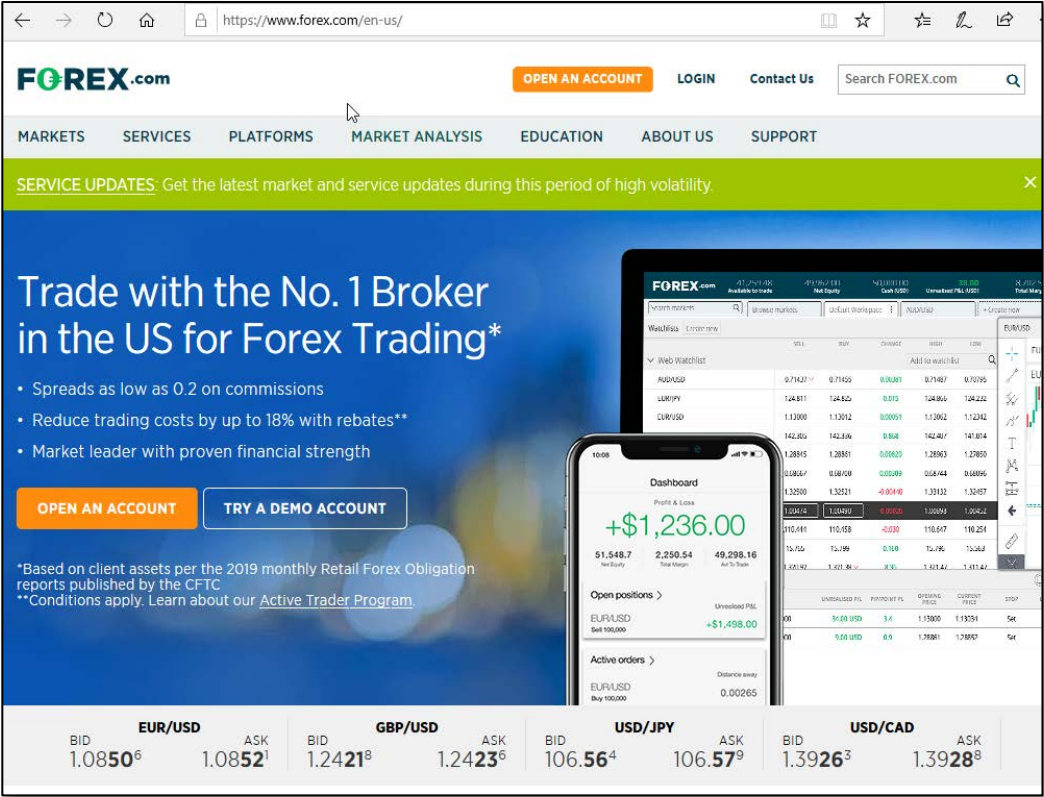
yield curve, 25

**Z**

zero-sum game, 353

Exhibit D

U.S. Patent 7,146,336

Claim Language	Gain Systems
<p><b>Claim 1</b></p>	
<p>1. A system for trading currencies over a computer network, comprising:</p>	<p>Defendants' systems practice this claim. The infringing "systems" include the web-based platform and back end servers, desktop clients, and mobile apps. <i>See Forex.com:</i></p>  <p>The screenshot shows the FOREX.com website interface. At the top, there is a navigation bar with 'FOREX.com' logo, 'OPEN AN ACCOUNT', 'LOGIN', and 'Contact Us' buttons. Below this is a menu with 'MARKETS', 'SERVICES', 'PLATFORMS', 'MARKET ANALYSIS', 'EDUCATION', 'ABOUT US', and 'SUPPORT'. A green banner reads 'SERVICE UPDATES: Get the latest market and service updates during this period of high volatility'. The main content area features a large blue banner with the text 'Trade with the No. 1 Broker in the US for Forex Trading*'. Below this banner are three bullet points: 'Spreads as low as 0.2 on commissions', 'Reduce trading costs by up to 18% with rebates**', and 'Market leader with proven financial strength'. There are two buttons: 'OPEN AN ACCOUNT' and 'TRY A DEMO ACCOUNT'. At the bottom of the main content area, there is a table of currency rates for EUR/USD, GBP/USD, USD/JPY, and USD/CAD. On the right side, there is a 'Web Watchlist' table with columns for currency pairs, bid, ask, and change. Below the watchlist is a 'Dashboard' section showing 'Profit &amp; Loss' of +\$1,236.00, 'Net Equity' of 51,548.7, 'Total Margin' of 2,250.54, and 'Av. To Trade' of 49,298.16. There are also sections for 'Open positions' and 'Active orders'.</p>

Oanda Corp. v. GAIN Capital Holdings, Inc.;  
GAIN Capital Group, LLC.

Exhibit D

(a) a server front-end in communication with said computer network;

See preamble and other screenshots. See also:

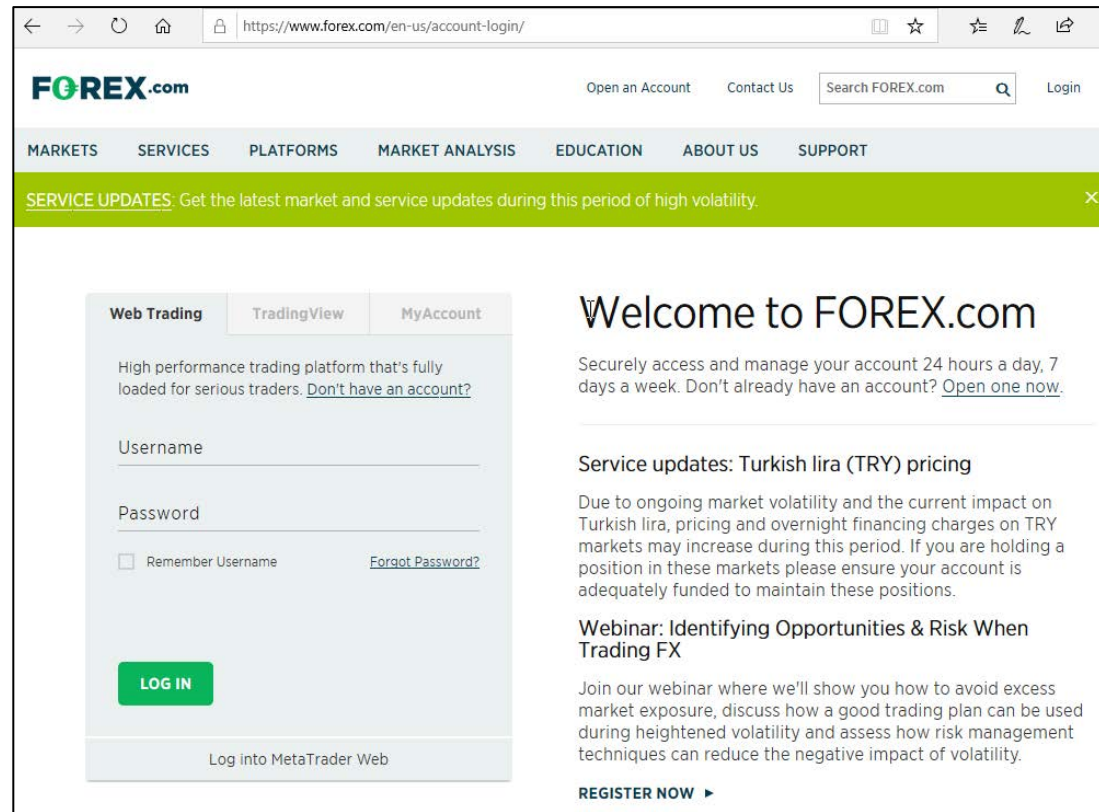
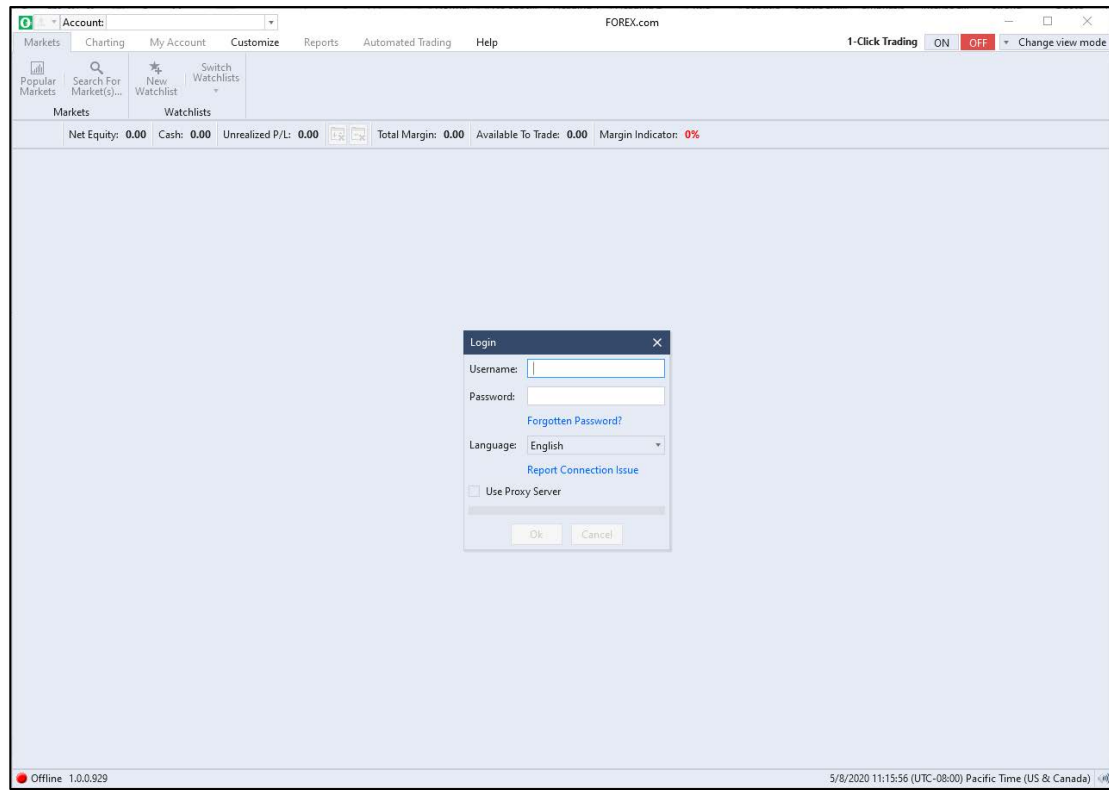


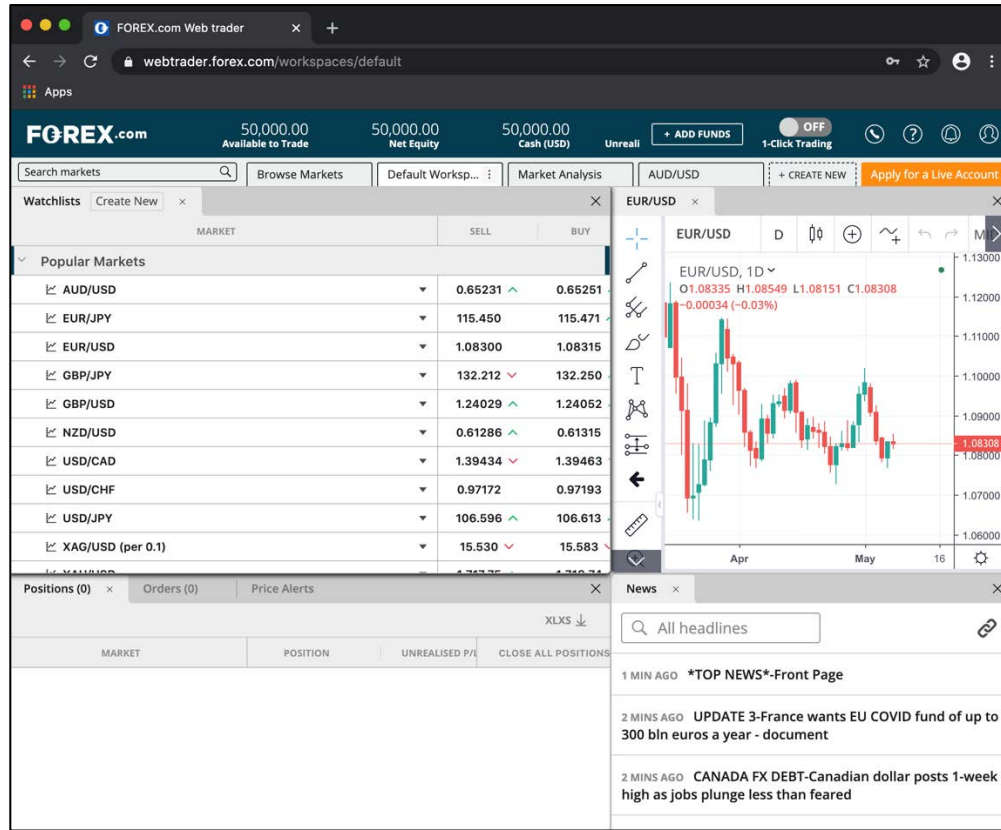
Exhibit D



(b) a database;

See:

Exhibit D

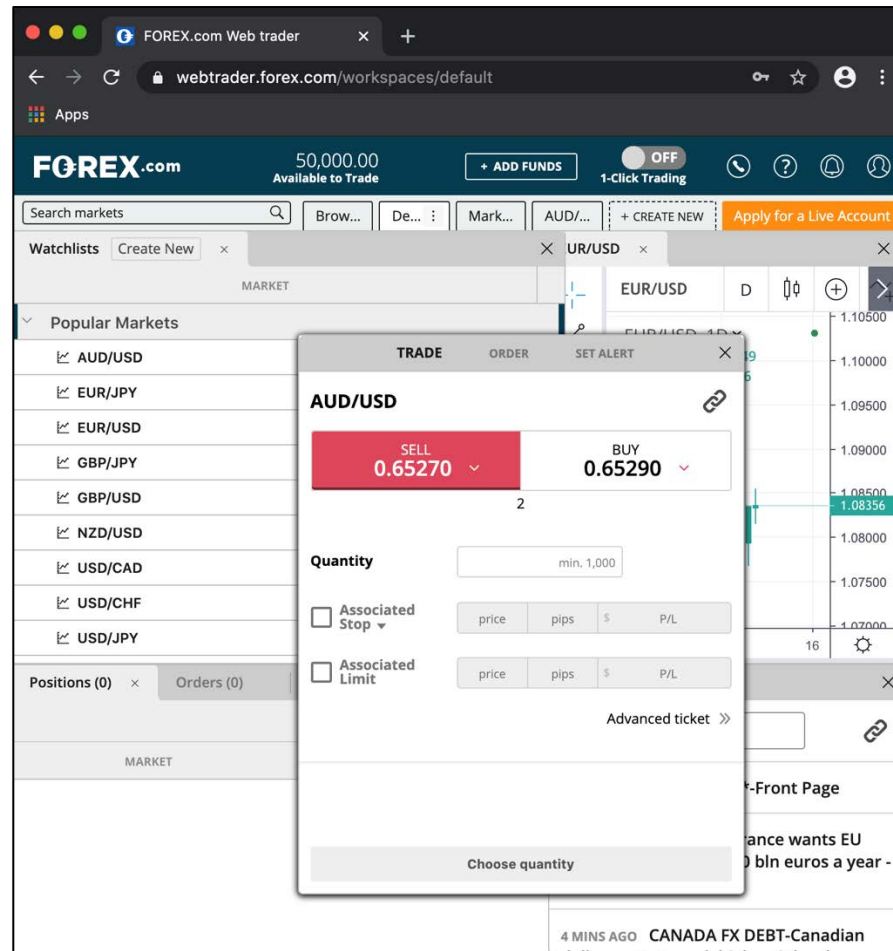


(c) a transaction server in communication with said server front-end and with said database;

See preamble and other screenshots. See also:

*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*

Exhibit D

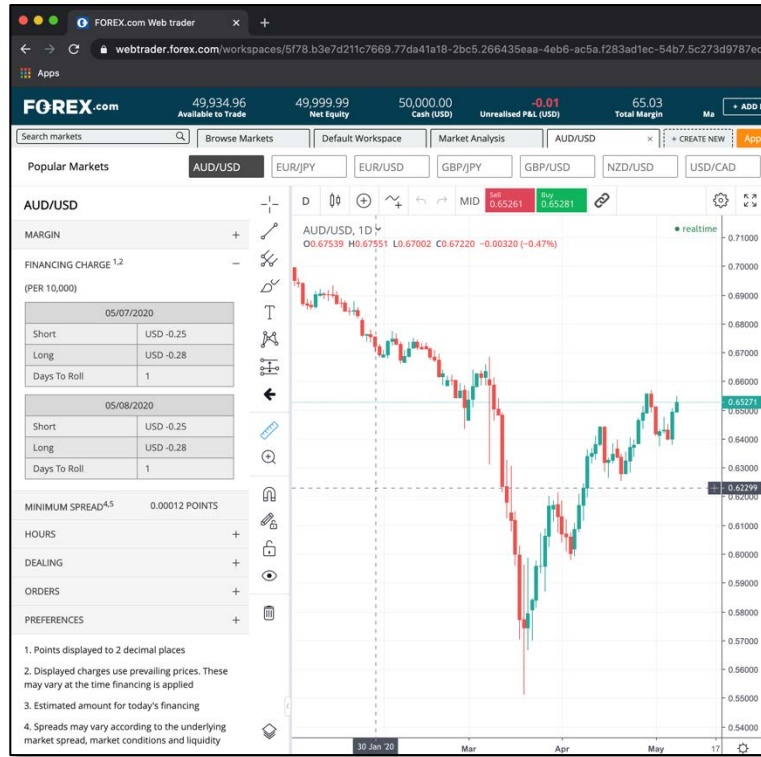


(d) a rate server in communication with said server front-end; and

See preamble and other screenshots. See also:

*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*

Exhibit D



(e) a pricing engine in communication with said rate server; and further comprising an interest rate manager in communication with said transaction server and said database, wherein said interest rate manager is operative to calculate, pay out, and collect interest on a tick-by-tick basis.

See preamble and other screenshots.

Exhibit D

flexible account solutions  
When you open an account with us, you're in control of your pricing.

ACCOUNT TYPE	Standard	Commission	Direct Market Access
DESCRIPTION	For traditional currency traders, costs to trade is bid/ask spread	For currency traders seeking consistently tight spreads - EUR/USD as low as 0.2 - with fixed commissions	For high-volume and frequent traders who demand a real-time view of market action with Level 2 pricing
FX SPREADS	Variable spreads, EUR/USD as low as 1.0. <a href="#">View spreads</a>	Raw spreads, EUR/USD as low as 0.2. <a href="#">View spreads</a>	Variable spreads as low as 0.1 on most major pairs. <a href="#">View spreads</a>
COMMISSIONS	None	Fixed \$5 per 100K traded	Variable starting at \$60 per million traded
VOLUME DISCOUNTS	Cash rebates of up to \$9 per million traded. <a href="#">Learn more</a>	Cash rebates of up to \$9 per million traded. <a href="#">Learn more</a>	Discounted commissions as low as \$20 per million traded. <a href="#">Learn more</a>
INTEREST PAID	Earn 1.5 APY interest on your average daily available margin balance up to \$50K, and 0.5 APY on \$50K-\$100K*		
TRADING PLATFORMS	FOREX.com & MetaTrader	FOREX.com	FOREX.com
ACCOUNT APPLICATIONS	<a href="#">OPEN A STANDARD ACCOUNT ▶</a> <a href="#">Learn more</a>	<a href="#">OPEN A COMMISSION ACCOUNT ▶</a> <a href="#">Learn more</a>	<a href="#">OPEN A DMA ACCOUNT ▶</a> <a href="#">Learn more</a>

Date	Details	Closing Serial No	Opening Serial No	Action	Is Close By	Quantity	Remaining Quantity	Opening Price	Closing Price	Account Ccy	Spread Cost	Realised P&L	Convert Rate	Trade Ccy	Trade P&L
	USD AUD/USD			Buy	No	2000.0000	2000.0000	0.642230		USD	-0.16			USD	
	USD AUD/USD			Buy	No	2000.0000	2000.0000	0.64149		USD	-0.19			USD	
	USD EUR/USD			Buy	No	2000.0000	2000.0000	1.08007		USD	-0.13			USD	
	USD USD/CHF FINANCING					0.0000	0.0000	0.98856		USD		-0.1900	1.02138298	CHF	-0.1840
	USD EUR/USD FINANCING					0.0000	0.0000	1.11154		USD		-0.1900			
	USD AUD/USD FINANCING					0.0000	0.0000	0.68985		USD		-0.0800			
	USD USD/CHF			Buy	No	2000.0000	0.0000	0.98856	0.98402	USD	-0.93	5.1400	1.01114803	CHF	3.0800
	USD EUR/USD			Sell	No	2000.0000	0.0000	1.11154	1.11597	USD	-0.64	8.8600			



Exhibit D

MARKET	POSITION	UNREALIZED PL	POINT PL	OPENING PRICE	CURRENT PRICE	STOP	LIMIT	MARGIN	AUTO ROLLER	DATE CHANGED	CLOSE ALL POSITIONS	
AUD/USD	Buy	4,000	2.46 USD	6.1	0.64140	0.64201	0.64100	Set	77.04		Close All	
EUR/USD	Buy	2,000	2.26 USD	15.3	1.08007	1.08120	1.07977	1.08157	43.25	—	06/05/2020 11:22:18	Close

<b>Claim 2</b>	
2. A system for trading currencies over a computer network, comprising:	See preamble to Claim 1.
(a) a server front-end in communication with said computer network;	See Claim 1(a).
(b) a database;	See Claim 1(b).
(c) a transaction server in communication with said server front-end and with said database;	See Claim 1(c).
(d) a rate server in communication with said server front-end; and	See Claim 1(d).
(e) a pricing engine in communication with said rate server; and further comprising a trade manager in communication with said transaction server and said database, wherein said trade manager comprises a stop-loss daemon that (a) continuously checks	See Claim 1(e). See also:

Exhibit D

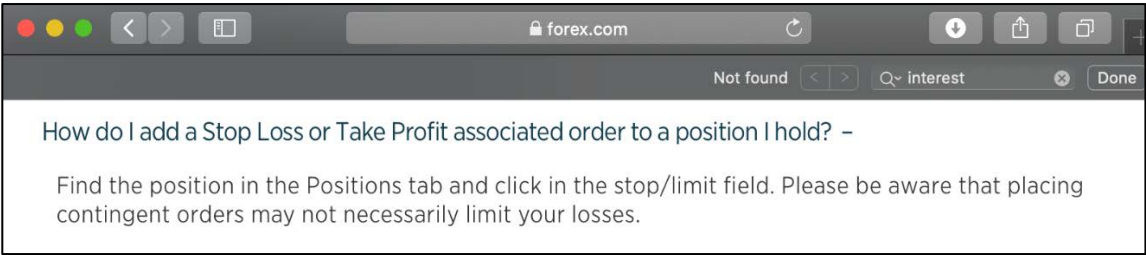
<p>whether stop-loss orders should be executed, and (b) if a stop-loss order should be executed, executes it through said transaction server.</p>	
<p><b>Claim 3</b></p>	
<p>3. A system for trading currencies over a computer network, comprising:</p>	<p>See preamble to Claim 1.</p>
<p>(a) a server front-end in communication with said computer network;</p>	<p>See Claim 1(a).</p>
<p>(b) a database;</p>	<p>See Claim 1(b).</p>
<p>(c) a transaction server in communication with said server front-end and with said database</p>	<p>See Claim 1(c).</p>
<p>(d) a rate server in communication with said server front-end; and</p>	<p>See Claim 1(d).</p>
<p>(e) a pricing engine in communication with said rate server; and further comprising a trade manager in communication with said transaction server and said database, wherein said trade manager comprises a take-profit</p>	<p>See Claim 1(e). <i>See also:</i></p>

Exhibit D

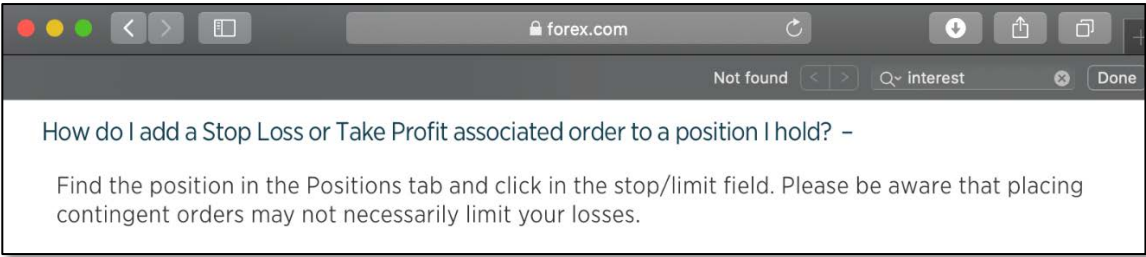
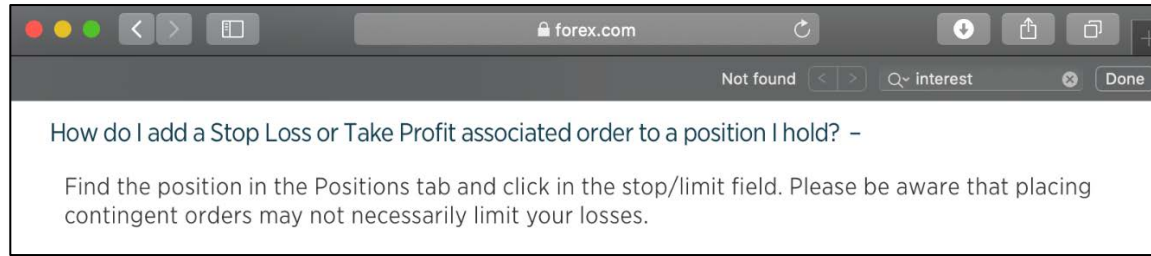
<p>daemon that (a) continuously checks whether take-profit orders should be executed, and (b) if a take-profit order should be executed, executes it through said transaction server.</p>	
<p><b>Claim 4</b></p>	
<p>4. A system for trading currencies over a computer network, comprising:</p>	<p>See preamble to Claim 1.</p>
<p>(a) a server front-end in communication with said computer network;</p>	<p>See Claim 1(a).</p>
<p>(b) a database;</p>	<p>See Claim 1(b).</p>
<p>(c) a transaction server in communication with said server front-end and with said database;</p>	<p>See Claim 1(c).</p>
<p>(d) a rate server in communication with said server front-end; and</p>	<p>See Claim 1(d).</p>
<p>(e) a pricing engine in communication with said rate server; and further comprising a trade manager in communication with said transaction server and said database, wherein said trade manager comprises a limit-order daemon that (a) continuously checks</p>	<p>See Claim 1(e). See also:</p>

Exhibit D

whether limit orders should be executed, and (b) if a limit order should be executed, executes it through said transaction server.



MARKET	ORDER	ORDER PRICE	CURRENT PRICE	DISTANCE AWAY	LIMIT	STOP	ORDER TYPE	DATE CHANGED		
EUR/USD	Sell	2,000	1.10000	1.08082	0.01918	Set	Set	GTC	06/05/2020 11:20:00	Delete
AUD/USD	Buy	2,000	0.64000	0.64211	0.00211	Set	Set	GTC	06/05/2020 12:04:41	Delete
AUD/USD	Buy	2,000	0.64000	0.64211	0.00211	Set	Set	GTC	06/05/2020 12:05:22	Delete
USD/CAD	Buy	1,000	1.41200	1.41274	0.00074	Set	Set	GTC	06/05/2020 12:56:21	Delete

Market	Market Expiry	Buy/Sell	Order Type	Quantity	Order Price	Current Price	Distance Away	Stop	Limit	Expiry	Order ID	Date/Time	Basis	ParentID	Expiry
EUR/USD		Sell	Limit	2,000	1.1	1.08059	0.01941			Good Till Cancelled	732593048	2020-05-06 11:20:00 AM	Single		
AUD/USD		Buy	Limit	2,000	0.64	0.64187	0.00187			Good Till Cancelled	732606240	2020-05-06 12:04:41 PM	Single		
AUD/USD		Buy	Limit	2,000	0.64	0.64187	0.00187			Good Till Cancelled	732606441	2020-05-06 12:05:22 PM	Single		
AUD/USD		Sell	Associated Stop	4,000	0.641	0.64166	0.00066			Good Till Cancelled	732608284	2020-05-06 12:13:09 PM	Single		
EUR/USD		Sell	Associated Stop	2,000	1.07977	1.08059	0.00082			Good Till Cancelled	732614372	2020-05-06 12:47:57 PM	Single		
EUR/USD		Sell	Associated Li...	2,000	1.08157	1.08059	0.00098			Good Till Cancelled	732615208	2020-05-06 12:47:57 PM	Single		
USD/CAD		Buy	Limit	1,000	1.412	1.41346	0.00146			Good Till Cancelled	732616602	2020-05-06 12:56:22 PM	Single		

**Claim 5**

5. A system for trading currencies over a computer network, comprising:

See preamble to Claim 1.

(a) a server front-end in communication with said computer network;

See Claim 1(a).

Exhibit D

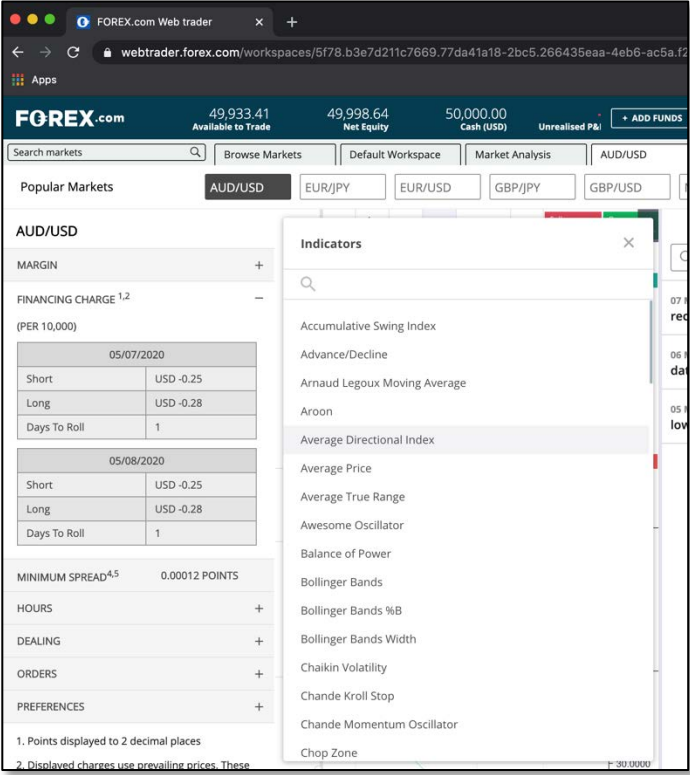
<p>(b) a database;</p>	<p>See Claim 1(b).</p>
<p>(c) a transaction server in communication with said server front-end and with said database;</p>	<p>See Claim 1(c).</p>
<p>(d) a rate server in communication with said server front-end; and</p>	<p>See Claim 1(d).</p>
<p>(e) a pricing engine in communication with said rate server, wherein said pricing engine is operable to compute currency exchange rates based on: (a) data obtained from external rate sources; and (b) market directional movement and volatility.</p>	<p>See Claim 1(e). <i>See also:</i></p>  <p>The screenshot shows the FOREX.com Web trader interface for the AUD/USD market. The top navigation bar includes 'FOREX.com', account balances (49,933.41 Available to Trade, 49,998.64 Net Equity, 50,000.00 Cash (USD)), and an 'ADD FUNDS' button. The main content area displays 'AUD/USD' with a 'Popular Markets' section and a 'MARGIN' table. The 'MARGIN' table shows financing charges for 05/07/2020 and 05/08/2020. A 'MINIMUM SPREAD' of 0.00012 POINTS is also displayed. An 'Indicators' dropdown menu is open, listing various technical indicators such as Accumulative Swing Index, Average Directional Index, and Bollinger Bands.</p>

Exhibit D

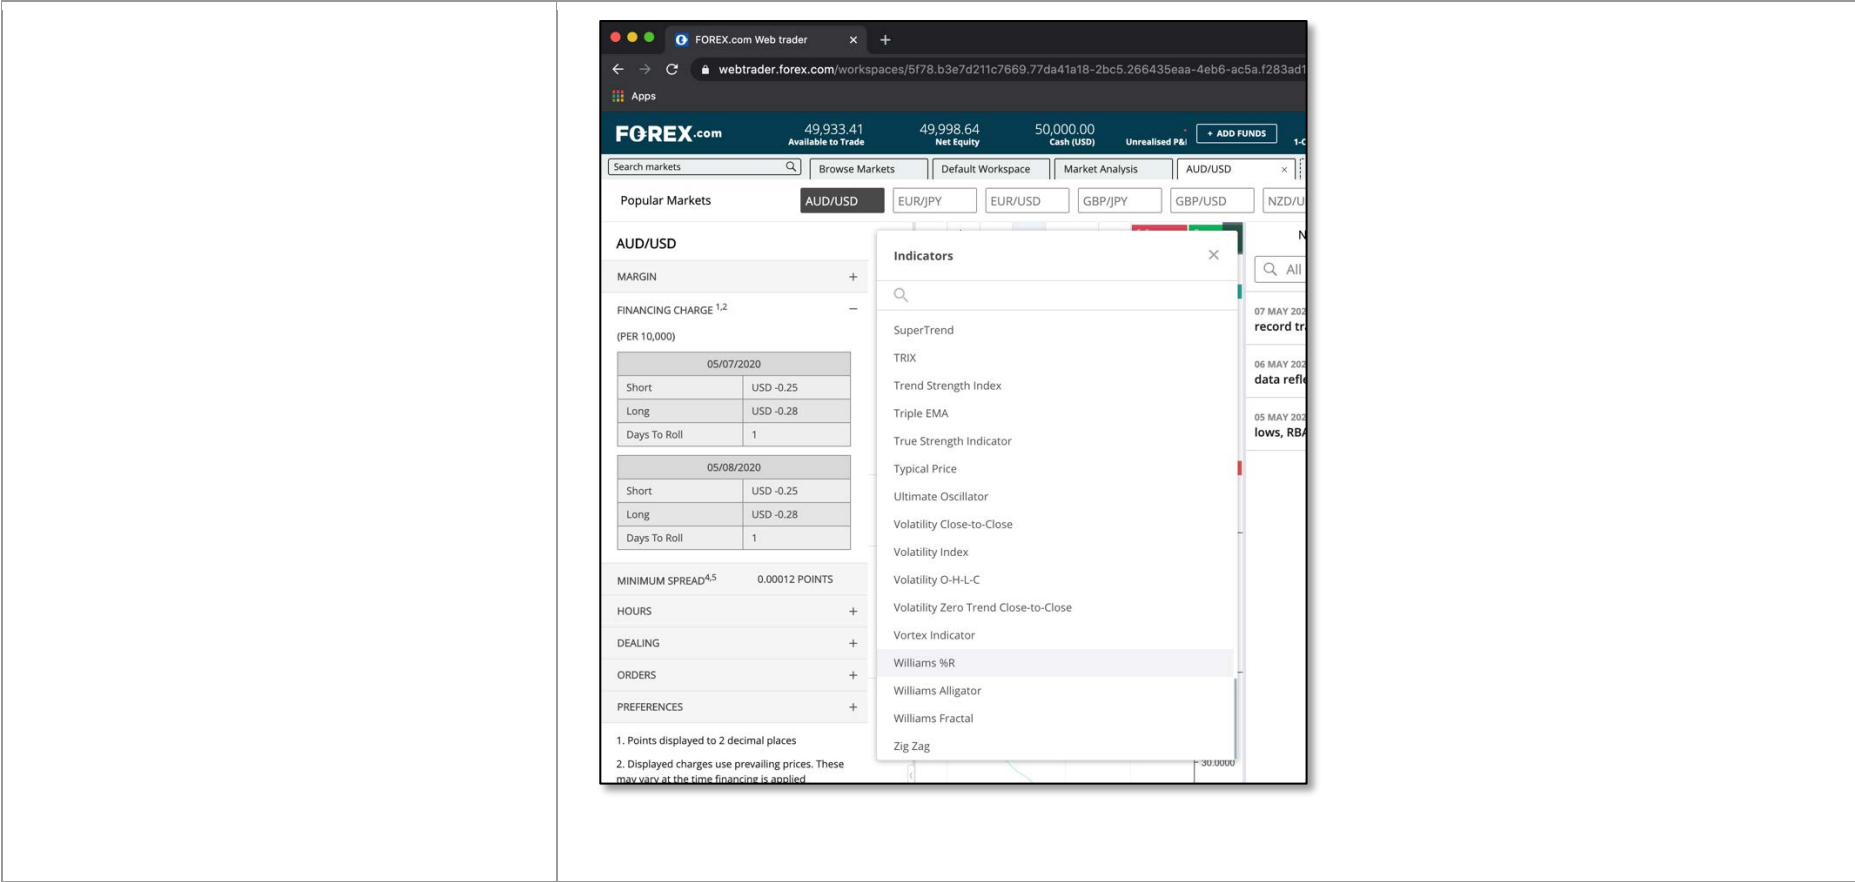
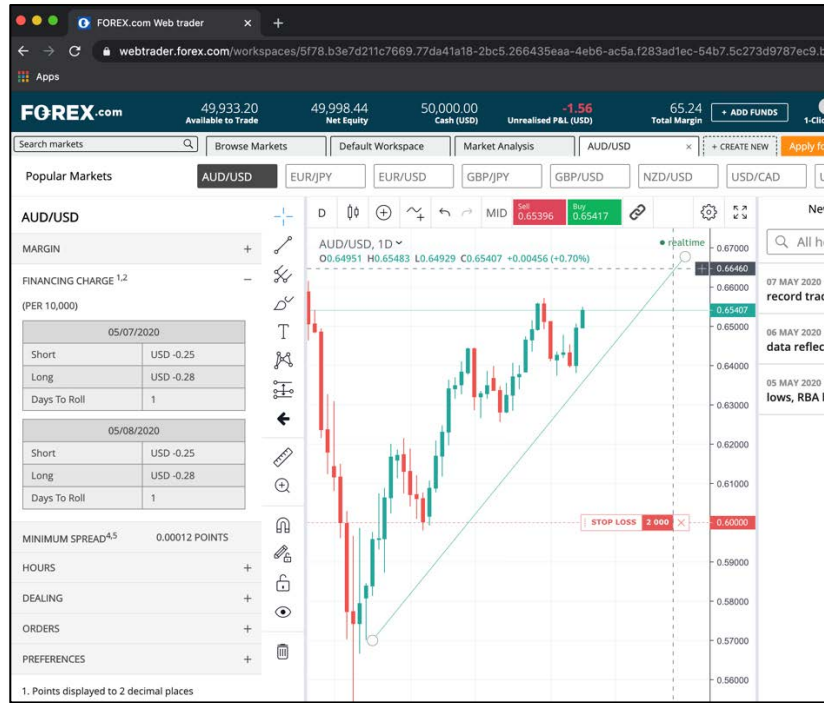


Exhibit D



**Claim 6**

6. A system as in claim 5, wherein said pricing engine is further operable to compute currency exchange rates based on positions held by said system.

See Claim 5. See also:

*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*

Exhibit D

<p><b>Claim 7</b></p>	
<p>7. A system for trading currencies over a computer network, comprising</p>	<p>See preamble to Claim 1.</p>



Exhibit D

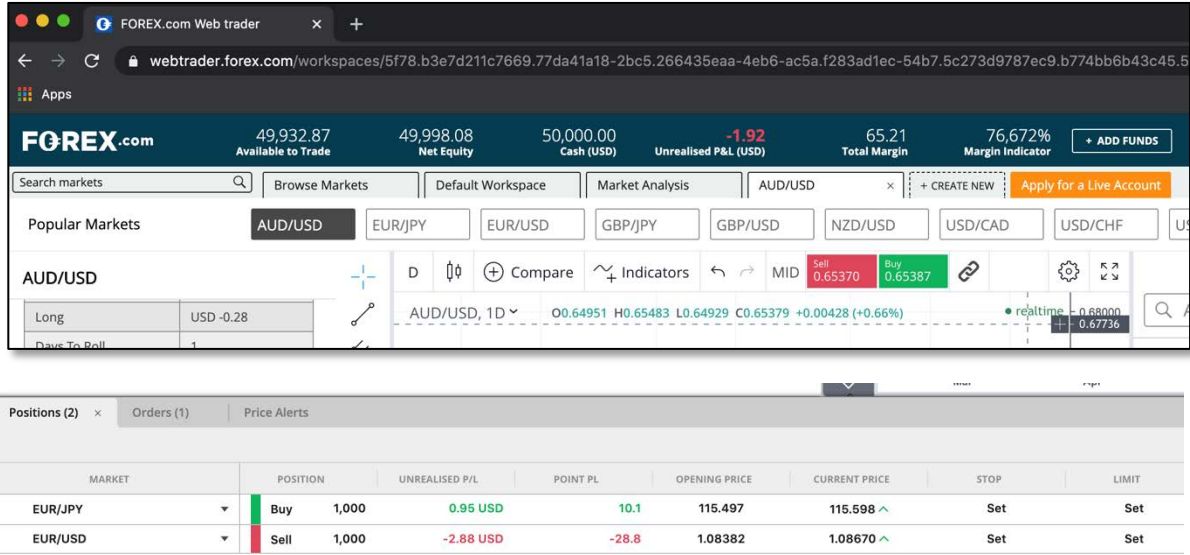
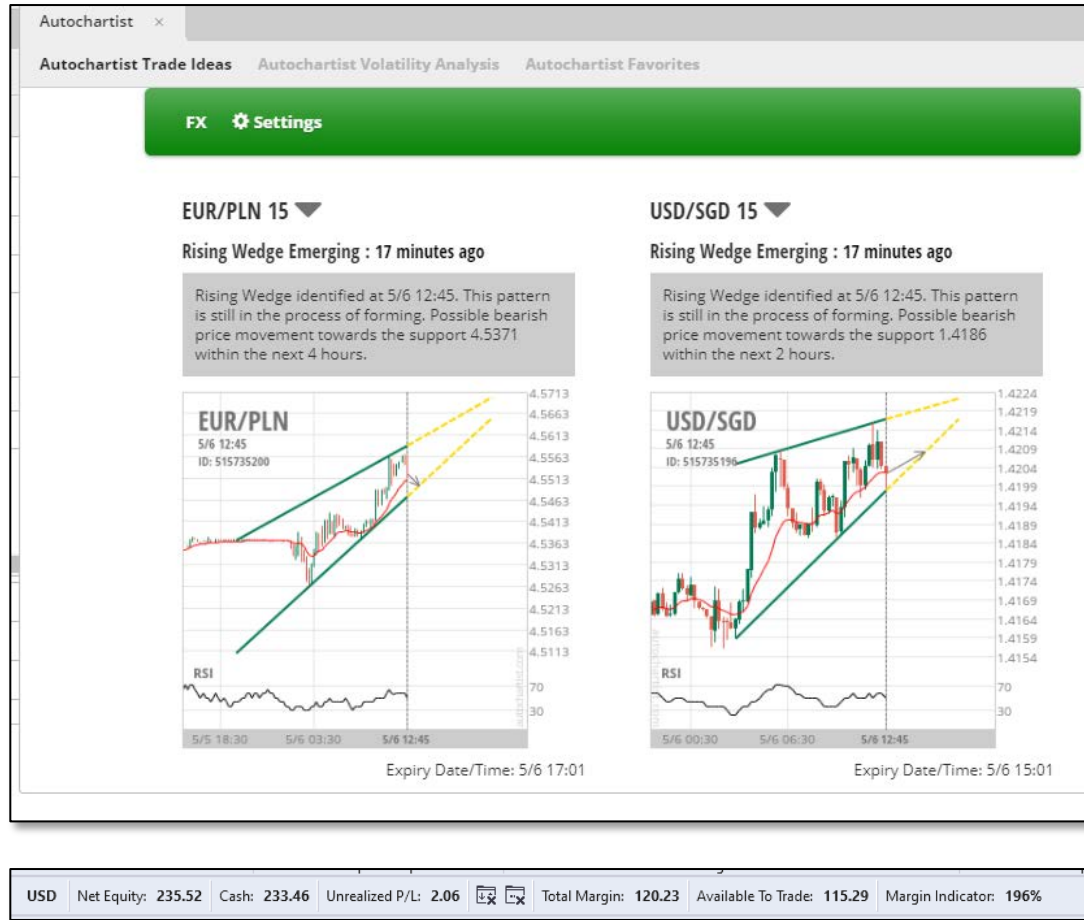
<p>(a) a server front-end in communication with said computer network;</p>	<p>See Claim 1(a).</p>																								
<p>(b) a database;</p>	<p>See Claim 1(b).</p>																								
<p>(c) a transaction server in communication with said server front-end and with said database;</p>	<p>See Claim 1(c).</p>																								
<p>(d) a rate server in communication with said server front-end; and</p>	<p>See Claim 1(d).</p>																								
<p>(e) a pricing engine in communication with said rate server; further comprising a hedging engine in communication with said transaction server, wherein said hedging engine is operable to perform at least two of the following calculations:                  (a) calculate a total amount of home currency appearing in all open positions;                  (b) calculate an out-of-equilibrium exposure; (c) calculate a new potential net exposure; (d) calculate an equilibrium position; (e) calculate boundaries of possible exposures; (f) calculate values for a pair of quoting functions; and (g) calculate an average price and an average spread.</p>	<p>See Claim 5. See also:</p>  <p>The screenshot shows the FOREX.com Web trader interface. At the top, account balances are displayed: Available to Trade (49,932.87), Net Equity (49,998.08), Cash (USD) (50,000.00), Unrealised P&amp;L (USD) (-1.92), Total Margin (65.21), and Margin Indicator (76,672%). Below this, there are tabs for 'Positions (2)', 'Orders (1)', and 'Price Alerts'. The 'Positions' tab is active, showing a table of open positions:</p> <table border="1"> <thead> <tr> <th>MARKET</th> <th>POSITION</th> <th>UNREALISED P/L</th> <th>POINT PL</th> <th>OPENING PRICE</th> <th>CURRENT PRICE</th> <th>STOP</th> <th>LIMIT</th> </tr> </thead> <tbody> <tr> <td>EUR/JPY</td> <td>Buy 1,000</td> <td>0.95 USD</td> <td>10.1</td> <td>115.497</td> <td>115.598</td> <td>Set</td> <td>Set</td> </tr> <tr> <td>EUR/USD</td> <td>Sell 1,000</td> <td>-2.88 USD</td> <td>-28.8</td> <td>1.08382</td> <td>1.08670</td> <td>Set</td> <td>Set</td> </tr> </tbody> </table>	MARKET	POSITION	UNREALISED P/L	POINT PL	OPENING PRICE	CURRENT PRICE	STOP	LIMIT	EUR/JPY	Buy 1,000	0.95 USD	10.1	115.497	115.598	Set	Set	EUR/USD	Sell 1,000	-2.88 USD	-28.8	1.08382	1.08670	Set	Set
MARKET	POSITION	UNREALISED P/L	POINT PL	OPENING PRICE	CURRENT PRICE	STOP	LIMIT																		
EUR/JPY	Buy 1,000	0.95 USD	10.1	115.497	115.598	Set	Set																		
EUR/USD	Sell 1,000	-2.88 USD	-28.8	1.08382	1.08670	Set	Set																		

Exhibit D



**Claim 8**

8. A system as in claim 6, wherein said positions are managed based on one or more trading models.

See Claim 6. See also:

Exhibit D

The screenshot displays the FOREX.com website's API Trading page. At the top, there is a navigation bar with the FOREX.com logo and a lock icon. Below the navigation bar is a green banner with the text "SERVICE UPDATES: Get the latest market and service updates during this period of high volatility." The main heading is "API Trading". The text below the heading reads: "Automate your trading by connecting your algo-trading strategies with our deep liquidity." and "Our REST API provides access to live streaming prices, trade execution, advanced order types, and access to over 80 of the world's most traded markets." A list of features includes: "Access to over 80 fx markets", "Execute trades and orders using trading systems and algos", and "Full developer resources". A green button labeled "OPEN AN ACCOUNT" is present, along with a link "Get started with a live MT4 account". A large graphic of three upward-pointing arrows (yellow, purple, and yellow) is on the right. The bottom section is titled "Range of REST API functionality" and contains four columns, each with an icon and a description: "Automated trading strategies" (with a refresh icon), "Integrated account management" (with a lock icon), "Charting and analysis" (with a bar chart icon), and "Easily compatible solution" (with a network icon).

FOREX.com

SERVICE UPDATES: Get the latest market and service updates during this period of high volatility.

## API Trading

Automate your trading by connecting your algo-trading strategies with our deep liquidity.

Our REST API provides access to live streaming prices, trade execution, advanced order types, and access to over 80 of the world's most traded markets.

- Access to over 80 fx markets
- Execute trades and orders using trading systems and algos
- Full developer resources

[OPEN AN ACCOUNT](#)

[Get started with a live MT4 account](#)

### Range of REST API functionality





Automated trading strategies	Integrated account management	Charting and analysis	Easily compatible solution
 Execute trades and a full range of orders against live streaming prices using your own algorithms or trading systems	 View your current active orders, account balance, available margin, open positions and historical trades in real-time.	 Strengthen your strategy with historical market data for deeper technical analysis	 Code against the API using any network accessible programming language from Perl-script, C++, Python or VB.NET.

Exhibit D

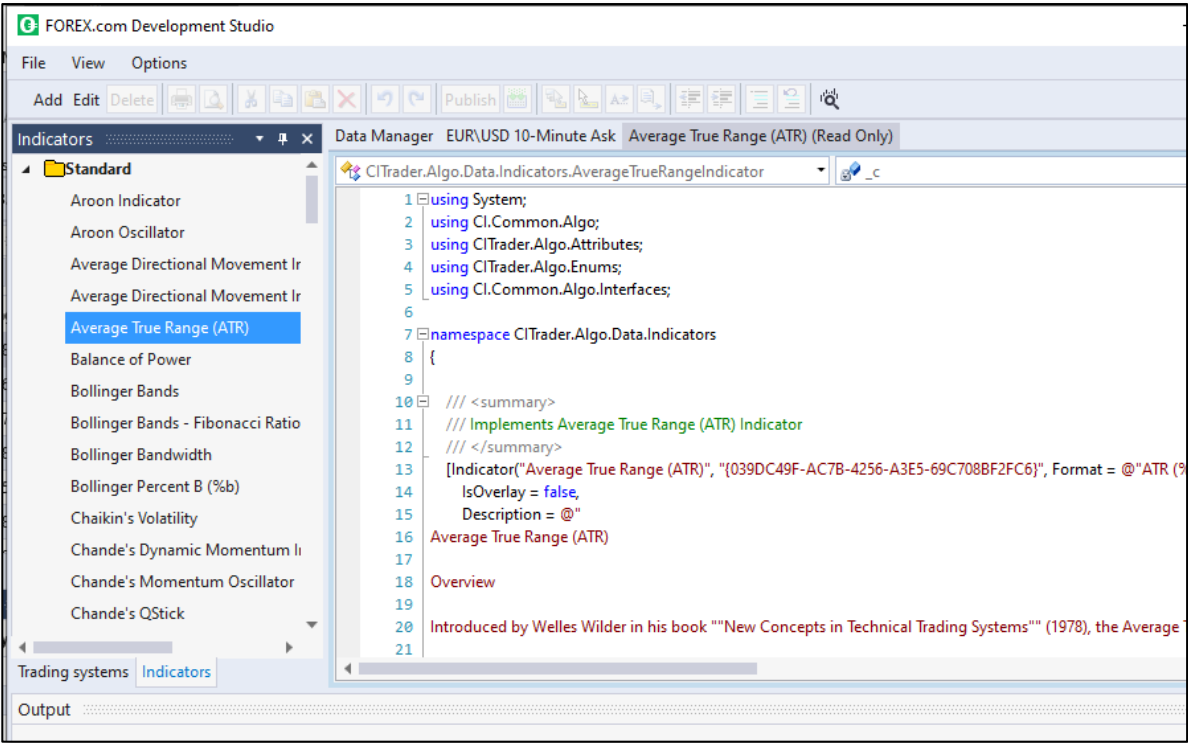
	 <p>The screenshot shows the FOREX.com Development Studio interface. On the left, a tree view under 'Indicators' has 'Average True Range (ATR)' selected. The main window displays the code for 'CITrader.Algo.Data.Indicators.AverageTrueRangeIndicator'. The code includes using statements for System, CI.Common.Algo, CITrader.Algo.Attributes, CITrader.Algo.Enums, and CI.Common.Algo.Interfaces. It defines a namespace 'CITrader.Algo.Data.Indicators' and implements the 'Average True Range (ATR)' indicator with a summary, format string, and description.</p> <pre> 1 using System; 2 using CI.Common.Algo; 3 using CITrader.Algo.Attributes; 4 using CITrader.Algo.Enums; 5 using CI.Common.Algo.Interfaces; 6 7 namespace CITrader.Algo.Data.Indicators 8 { 9 10     /// &lt;summary&gt; 11     /// Implements Average True Range (ATR) Indicator 12     /// &lt;/summary&gt; 13     [Indicator("Average True Range (ATR)", "{039DC49F-AC7B-4256-A3E5-69C708BF2FC6}", Format = @"ATR (% 14         IsOverlay = false, 15         Description = @" 16         Average True Range (ATR) 17 18     Overview 19 20     Introduced by Welles Wilder in his book ""New Concepts in Technical Trading Systems"" (1978), the Average 21 </pre>
<p><b>Claim 9</b></p>	<p>See Claim 8. See also:</p>
<p>9. A system as in claim 8, wherein at least one of said one or more trading models comprises: (a) a price collector component; (b) a price filter component; (c) a price database component; (d) a gearing calculator component; (e) a deal</p>	<p>See Claim 8. See also:</p>

Exhibit D

acceptor component; and (f) a book-keeper component.

So how do I find an ABCD pattern?

Each pattern has both a bullish and bearish version. Bullish patterns help identify higher probability opportunities to buy, or go "long." Bearish patterns help signal opportunities to "short," or sell.

Each turning point (A, B, C, and D) represents a significant high or significant low on a price chart. These points define three consecutive price swings, or trends, which make up each of the three pattern "legs." These are referred to as the AB leg, the BC leg, and the CD leg.



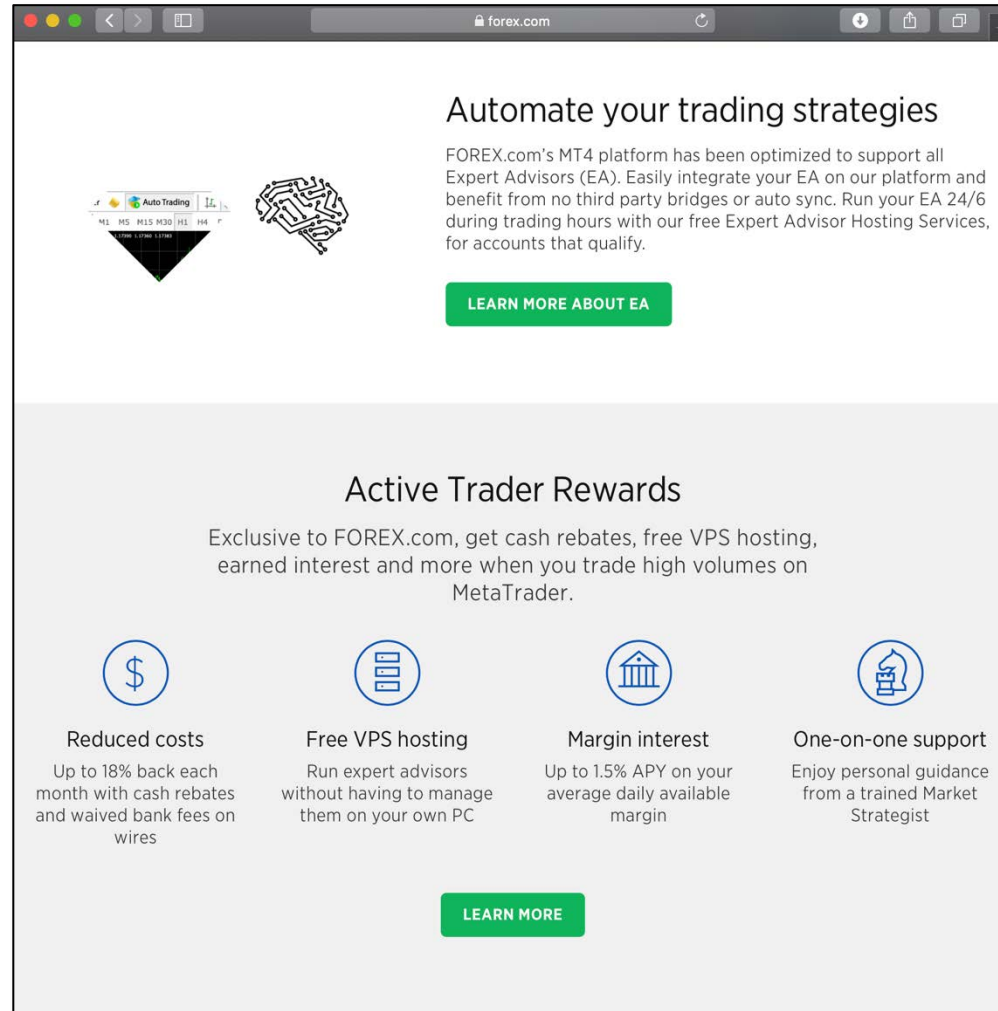
<https://www.forex.com/en-us/education/education-themes/technical-analysis/abcd-pattern/>  
 (accessed May 8, 2020).

Claim 10

Exhibit D

10. A system as in claim 8, wherein at least one of said one or more trading models comprises: (a) a price collector component; (b) a price filter component; (c) a price database component; (d) a gearing calculator component; (e) a deal acceptor component; (f) an opportunity catcher component; and (g) a book-keeper component

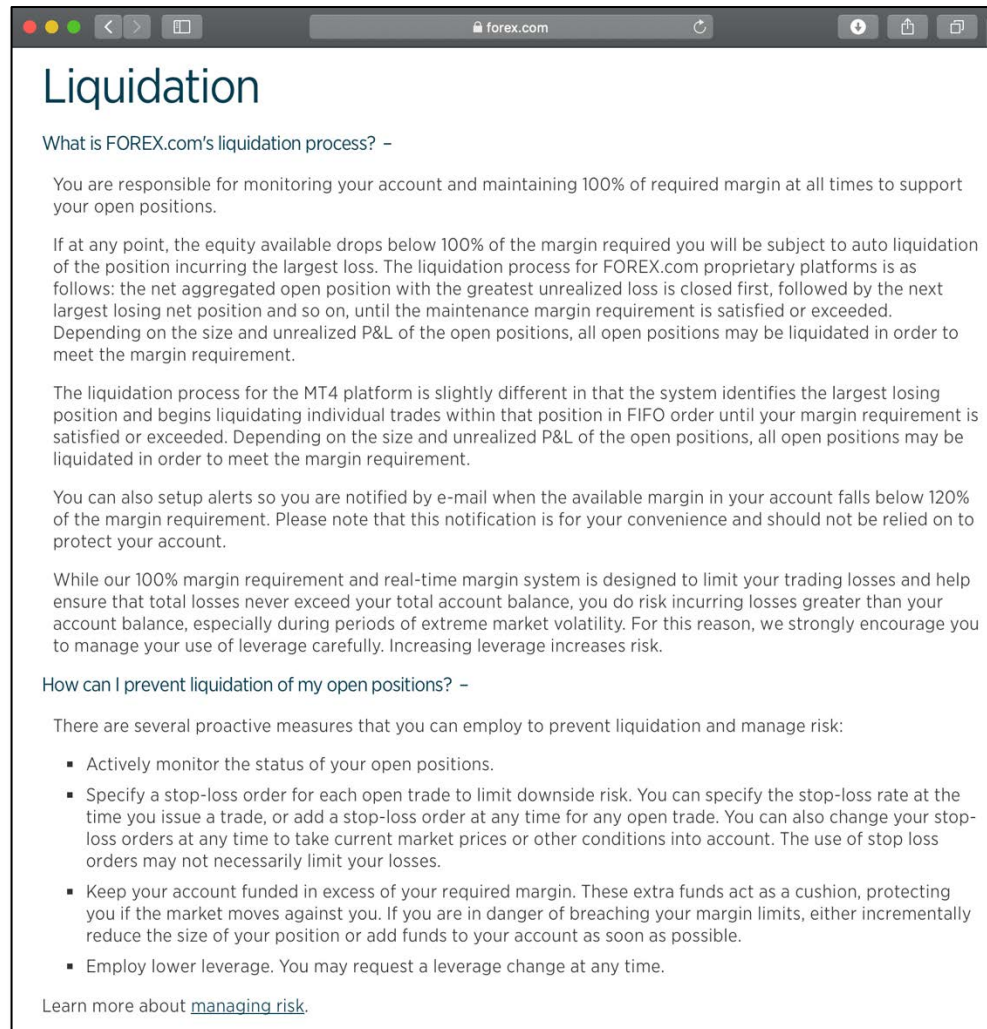
See Claim 9. See also:



## Exhibit D

Claim 11	
11. A system for trading currencies over a computer network, comprising:	See Claim 1.
(a) a server front-end in communication with said computer network;	See Claim 1(a).
(b) a database;	See Claim 1(b).
(c) a transaction server in communication with said server front-end and with said database;	See Claim 1(c).
(d) a rate server in communication with said server front-end; and	See Claim 1(d).
(e) a pricing engine in communication with said rate server; further comprising a margin control manager in communication with said transaction server and said database, wherein said margin control manager is operable to monitor on a tick-by-tick basis margin requirements of accounts and on said tick-by-tick basis liquidate holdings as needed to maintain specified margins.	See Claim 1(e). <i>See also:</i>

## Exhibit D



**Liquidation**

What is FOREX.com's liquidation process? –

You are responsible for monitoring your account and maintaining 100% of required margin at all times to support your open positions.

If at any point, the equity available drops below 100% of the margin required you will be subject to auto liquidation of the position incurring the largest loss. The liquidation process for FOREX.com proprietary platforms is as follows: the net aggregated open position with the greatest unrealized loss is closed first, followed by the next largest losing net position and so on, until the maintenance margin requirement is satisfied or exceeded. Depending on the size and unrealized P&L of the open positions, all open positions may be liquidated in order to meet the margin requirement.

The liquidation process for the MT4 platform is slightly different in that the system identifies the largest losing position and begins liquidating individual trades within that position in FIFO order until your margin requirement is satisfied or exceeded. Depending on the size and unrealized P&L of the open positions, all open positions may be liquidated in order to meet the margin requirement.

You can also setup alerts so you are notified by e-mail when the available margin in your account falls below 120% of the margin requirement. Please note that this notification is for your convenience and should not be relied on to protect your account.

While our 100% margin requirement and real-time margin system is designed to limit your trading losses and help ensure that total losses never exceed your total account balance, you do risk incurring losses greater than your account balance, especially during periods of extreme market volatility. For this reason, we strongly encourage you to manage your use of leverage carefully. Increasing leverage increases risk.

How can I prevent liquidation of my open positions? –

There are several proactive measures that you can employ to prevent liquidation and manage risk:

- Actively monitor the status of your open positions.
- Specify a stop-loss order for each open trade to limit downside risk. You can specify the stop-loss rate at the time you issue a trade, or add a stop-loss order at any time for any open trade. You can also change your stop-loss orders at any time to take current market prices or other conditions into account. The use of stop loss orders may not necessarily limit your losses.
- Keep your account funded in excess of your required margin. These extra funds act as a cushion, protecting you if the market moves against you. If you are in danger of breaching your margin limits, either incrementally reduce the size of your position or add funds to your account as soon as possible.
- Employ lower leverage. You may request a leverage change at any time.

Learn more about [managing risk](#).



# EXHIBIT E

Exhibit E

U.S. Patent 8,392,311

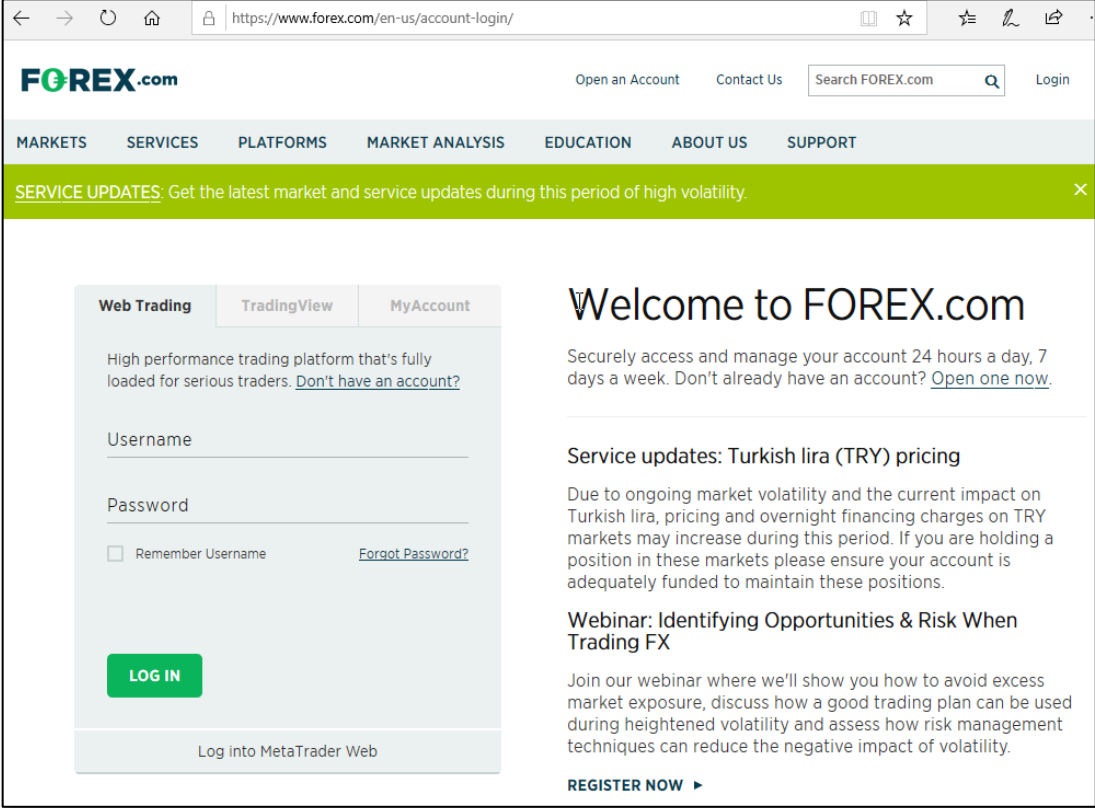
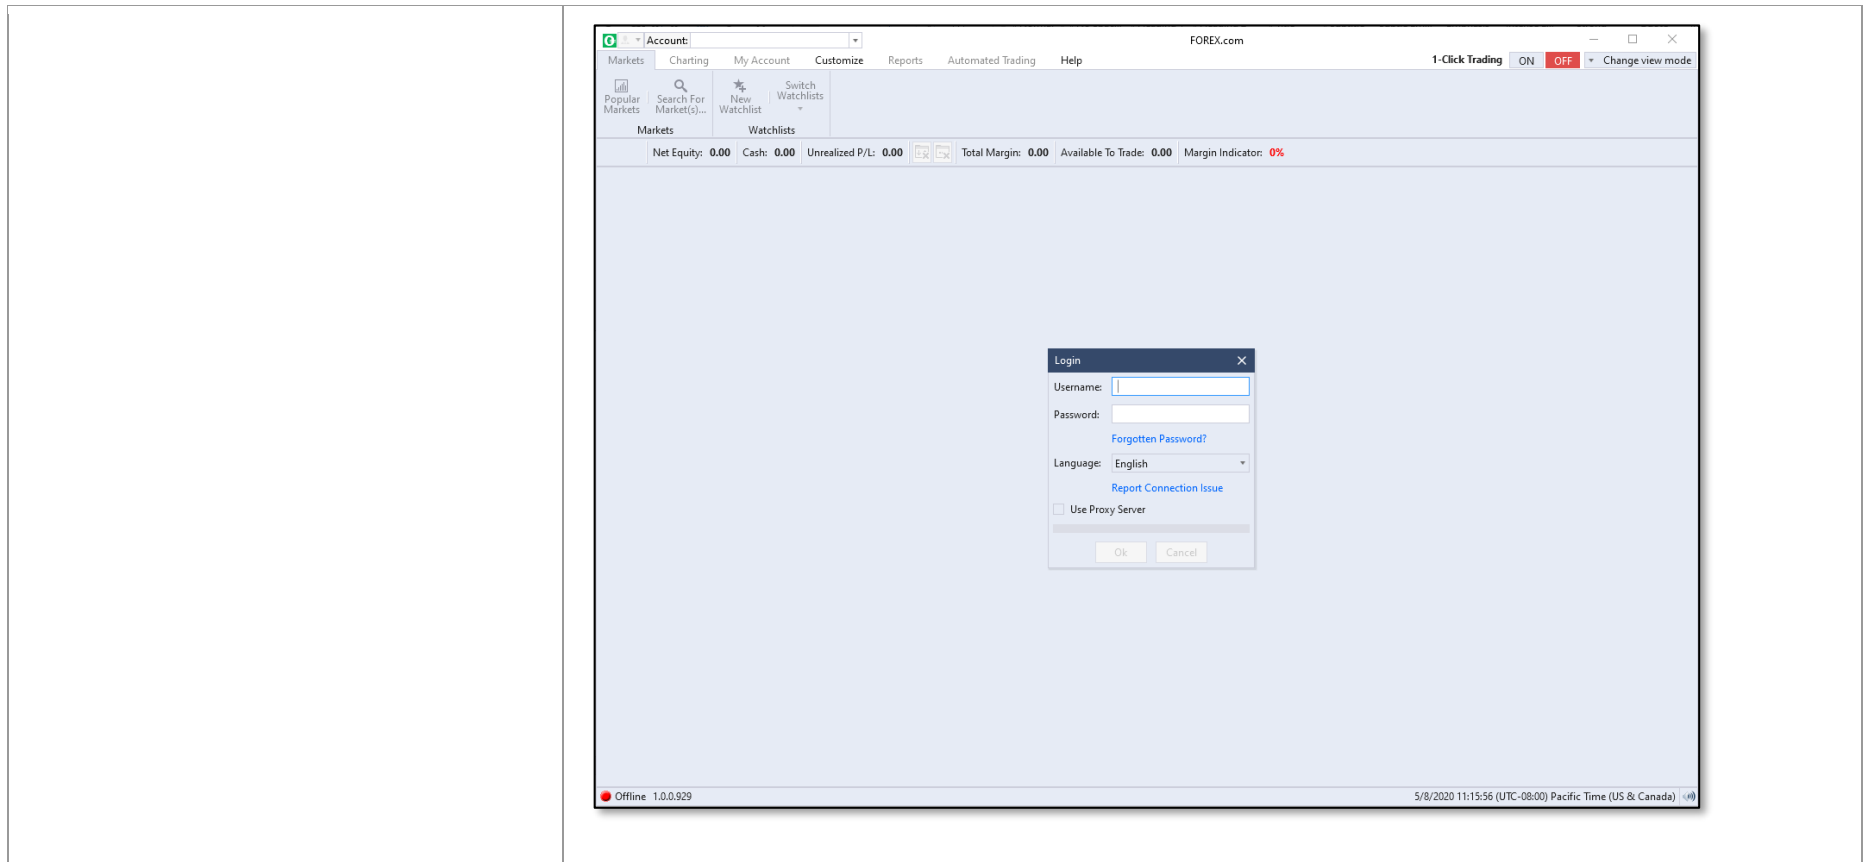
Claim Language	Gain Systems
<p align="center"><b>Claim 1</b></p>	
<p>1. A method of trading currencies over a computer network connecting a trading system server and at least one trading client system, comprising the steps of:</p>	<p>Defendants' currency trading systems practice this claim. See Forex.com:</p> 

Exhibit E



*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*

Exhibit E

(i) at the trading system server, determining and dynamically maintaining a plurality of current exchange rates, each current exchange rate relating to a pair of currencies and including a first price to buy a first currency of the pair with respect to a second currency of the pair and a second price to sell the first currency of the pair with respect to the second currency of the pair;

See:

The screenshot shows the FOREX.com interface with account balances at the top: Available to Trade (233.46), Net Equity (233.46), Cash (USD) (233.46), Unrealised P&L (USD) (0.00), Total Margin (0.00), and Margin Indicator (> 200%). The main content area is titled 'Popular' and 'POPULAR USA'. It features a table with columns for currency pairs, sell/buy prices, and percentage changes.

Filter markets	SELL	BUY	CHANGE	%CHANGE	
AUD/USD	0.64235	0.64257	-6.40000	-0.10	+
EUR/JPY	114.568	114.589	-92.600	-0.80	+
EUR/USD	1.08045	1.08061	-35.00000	-0.32	+
GBP/JPY	131.051	131.086	-147.700	-1.11	+
GBP/USD	1.23598 ^	1.23618	-74.70000	-0.60	+
NZD/USD	0.60267	0.60304	-22.70000	-0.38	+
USD/CAD	1.41363	1.41389	88.00000	0.63	+
USD/CHF	0.97455	0.97481	20.50000	0.21	+
USD/JPY	106.032	106.048	-53.600	-0.50	+
XAG/USD (per 0.1)	14.901	14.950	-0.920	-0.61	+
XAU/USD	1,686.13 v	1,687.14 ^	-19.42	-1.14	+

Exhibit E

Watchlist: Popular Markets										
Search For Market(s)...		Watchlist: Popular Markets								
Market	Bid (Sell)	Ask (Buy)	Spread	% Change		i		Trade Size	Order	
AUD/USD	0.64155 ▲	0.64176 ▲	2.1	-0.22%	☆	i		2,000	Order	
EUR/JPY	114.596 ▲	114.612 ▼	1.6	-0.78%	☆	i		3,000	Order	
EUR/USD	1.08036 ▼	1.08050 ▼	1.4	-0.33%	☆	i		2,000	Order	
GBP/JPY	130.975 ▲	131.011 ▲	3.6	-1.17%	☆	i		1,000	Order	
GBP/USD	1.23488 ▲	1.23508 ▲	2.0	-0.69%	☆	i		1,000	Order	
NZD/USD	0.60176 ▼	0.60209 ▼	3.3	-0.53%	☆	i		1,000	Order	
USD/CAD	1.41388 ▼	1.41410 ▲	2.2	0.64%	☆	i		10,000	Order	
USD/CHF	0.97449 ▼	0.97471 ▼	2.2	0.20%	☆	i		2,000	Order	
USD/JPY	106.066 ▼	106.078 ▼	1.2	-0.47%	☆	i		5,000	Order	
XAG/USD (per 0.1)	14.944 ▲	14.985 ▼	4.1	-0.35%	☆	i		5	Order	
XAU/USD	1,689.10 ▲	1,690.27 ▼	117	-0.96%	☆	i		1	Order	

Exhibit E

(ii) transmitting data from the trading system server to a trading client system, the transmitted data representing at least one current exchange rate at the time of the transmission;

See:




Oanda Corp. v. GAIN Capital Holdings, Inc.;  
GAIN Capital Group, LLC.

Exhibit E



*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*

Exhibit E



*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*



Exhibit E

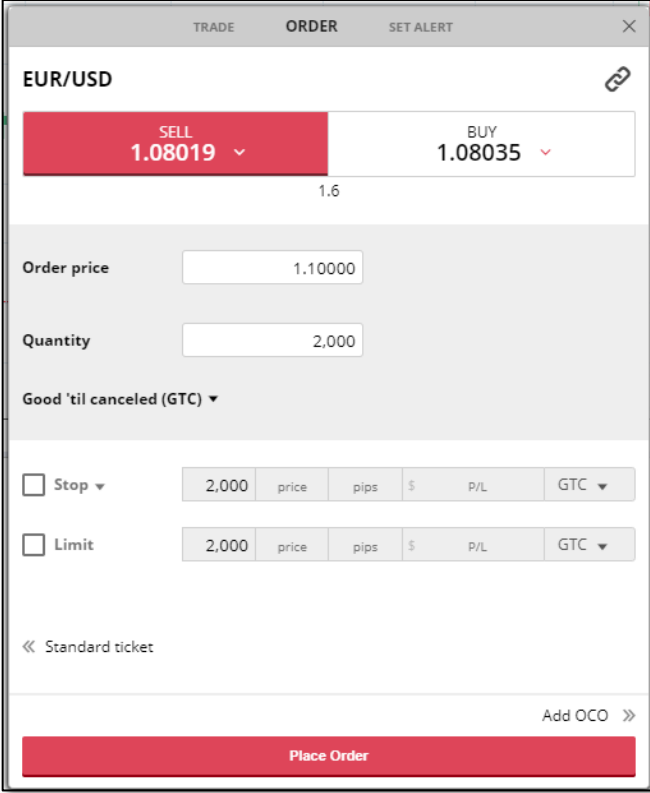
<p>(iii) at the trading client system, displaying the first and second prices for each received current exchange rate to a user;</p>	<p>See Claim 1(ii).</p>
<p>(iv) at the trading client system, accepting input from the user identifying a pair of currencies the user desires to trade, an amount of at least one currency of the pair desired to be traded and a requested trade price at which it is desired to effect the trade;</p>	<p>See:</p> 

Exhibit E

Create Entry Order

**EUR/USD**

1.08442 ▼      1.08459 ▼

Order Type: Entry ▼

Direction: Buy ▼

Price:       <= 1.08434      >= 1.08467

Quantity:       >= 1,000      <= 5,000,001

Good Till: Cancelled ▼

Stops/Limits Remove All

**Stop/Limit**

Stop       Limit

Trailing

Price:  <= 1.08434       >= 1.0845

PIPs:      

Estimated P/L:      

Quantity:      

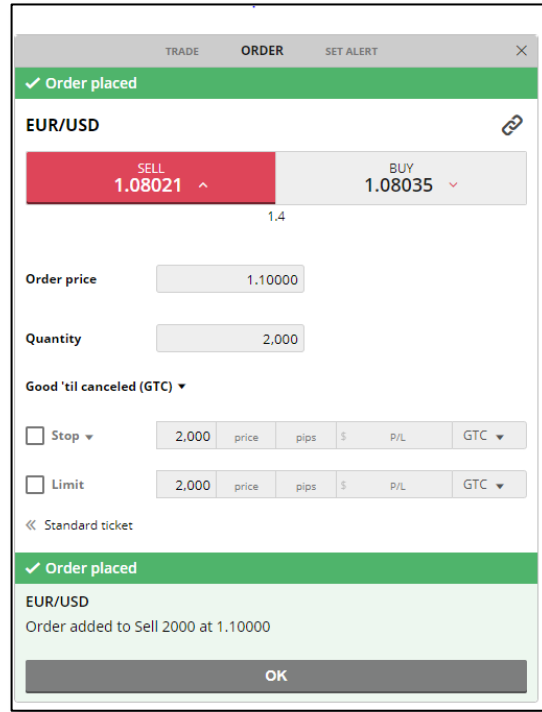
Good Till:

Ttl Qty: 0      Ttl Qty: 0

Exhibit E

(v) transmitting the accepted input from the trading client system to the trading system server;

See:



Positions (2)		Orders (1)	Price Alerts			
		ORDER	ORDER PRICE	CURRENT PRICE	DISTANCE AWAY	LI
AUD/USD	Buy	2,000	0.70000	0.65423	0.04577	Delete

Exhibit E

(vi) at the trading system server, comparing the requested trade price to the respective first price or second price of the corresponding current exchange rate at that time and, if the respective first price or second price of the corresponding current exchange rate at that time is equal to or better than the requested trade price, effecting the trade at the corresponding respective current exchange rate first price or second price and if the corresponding current exchange rate is worse than the requested trade price, refusing the trade; and

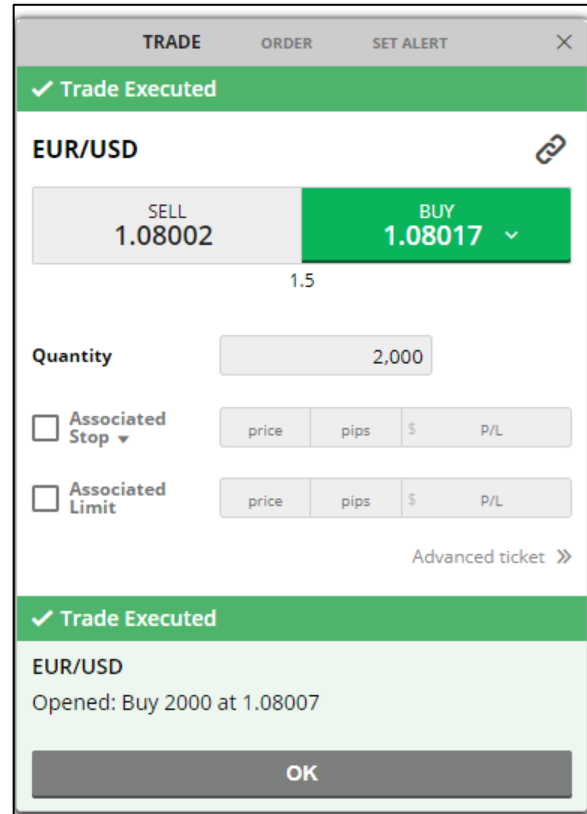


Exhibit E

TRADE
ORDER
SET ALERT
✕

### EUR/USD

SELL  
1.08001

BUY  
1.08017 ^

1.6

Quantity

Associated Stop ▾

<input style="width: 80%;" type="text" value="1.10"/>	-198.8	\$	79.52
---	--------	----	-------

Too high. Maximum price 1.07993

Associated Limit

price	pips	\$	P/L
-------	------	----	-----

Associated Quantity: 4,000  
Advanced ticket >>

Stop level too high

Time	News	Category
2020-05-06 12:02:30 PM	Italy's daily coronavirus death toll and new cases' tally both rise	Foreign Exchange
2020-05-06 12:02:23 PM	*TOP NEWS*-Central Banks & Global Economy	
2020-05-06 12:02:21 PM	*TOP NEWS*-Global Markets	
2020-05-06 12:00:32 PM	*TOP NEWS*-World News	
2020-05-06 11:58:18 AM	REFILE-UPDATE 2-Sterling hits 12-day low versus dollar	
2020-05-06 11:57:47 AM	UPDATE 1-Natixis swings to a loss, hikes provisions on coronavirus, frauds	
2020-05-06 11:55:51 AM	BUZZ-COMMENT-Yen gains counter upbeat risk, AUD/USD rally struggles	
2020-05-06 11:55:43 AM	WRAPUP 2-U.S. private payrolls dive by a record 20.2 million	

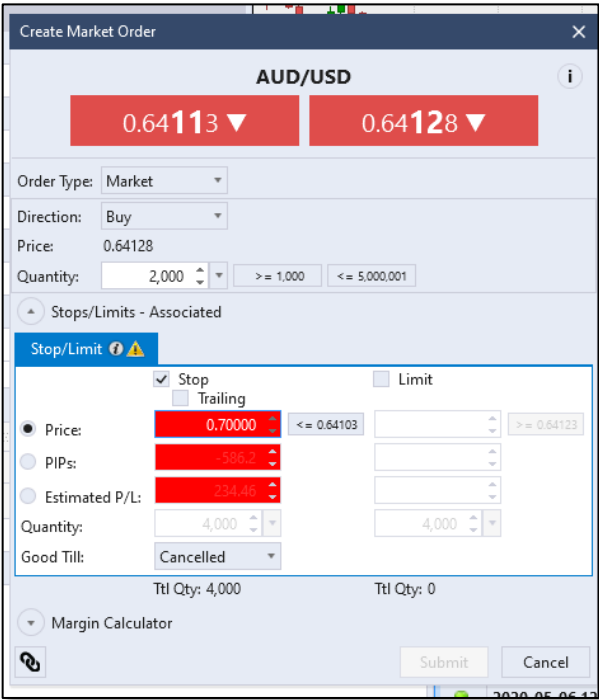
News | Economic Calendar

i

New trade: Market AUD/USD, BUY, 0.64149, Quantity 2,000, Order ID [REDACTED]

✕

Exhibit E



(vii) transmitting from the trading system server to the trading client system an indication of whether the trade was refused or transacted and, if transacted, an indication of the price the trade was transacted at.

See above claim elements.

Exhibit E

**Claim 2**

2. The method of claim 1 wherein the requested trade price is derived from a respective one of the first price or second price of the received current exchange rate and a user input limit value defining a maximum acceptable difference between the respective one of the first price or second price of the received current exchange rate received at the trading client system and the respective one of the first price or second price of the corresponding current exchange rate determined at the trading client system at which the trade can be effected.

See Claim 1. See also:

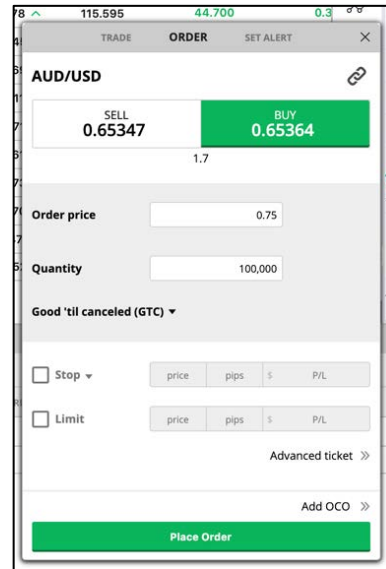
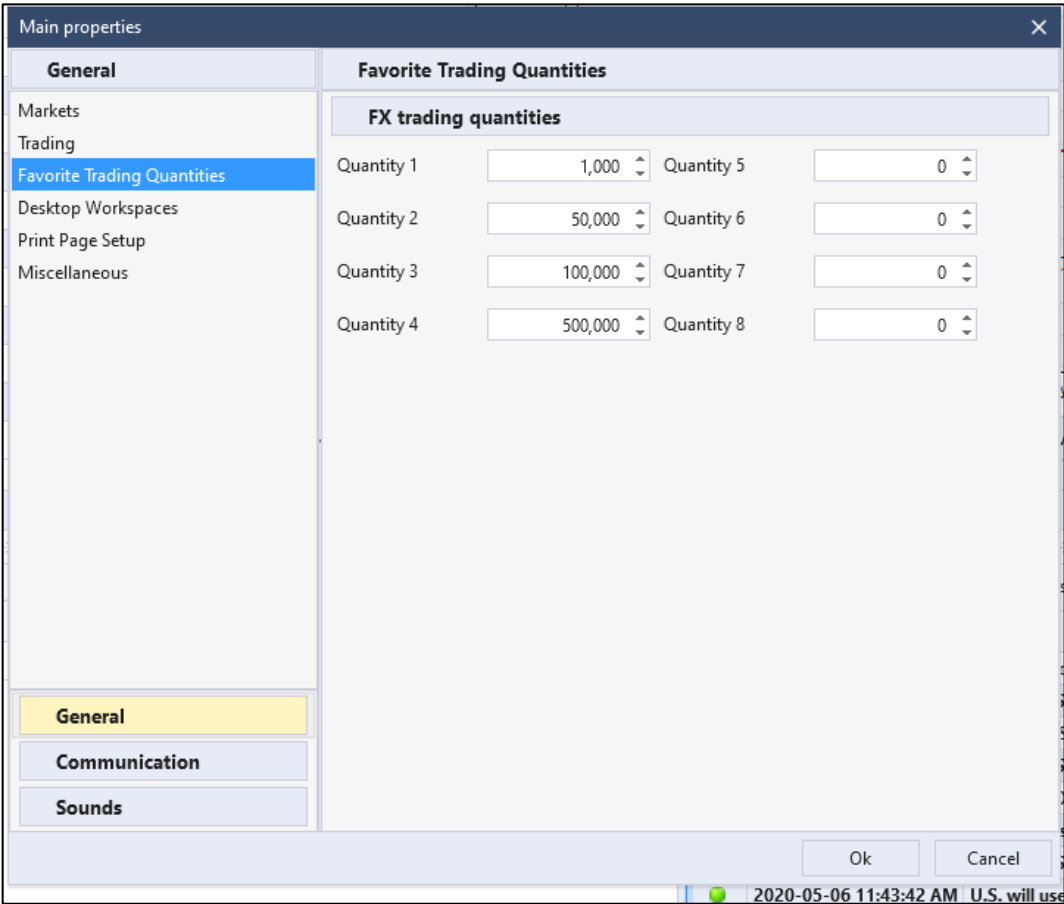


Exhibit E



See also: <https://www.forex.com/en-us/education/education-themes/forex-platform-tutorials/us-advanced-trading-platform-customization/> (accessed May 8, 2020)



## Exhibit E

<b>Claim 3</b>	
<p>3. The method of claim 2 wherein the user can input a first limit value to define a maximum acceptable difference between the first price of the current exchange rate received at the trading client system and the first price of the corresponding current exchange rate determined at the trading client system and can input a second limit value to define a maximum acceptable difference between the second price of the current exchange rate received at the trading system and the second price of the corresponding current exchange rate determined at the trading client system and the requested trade price is derived from the first price or second price of the current exchange rate received at the trading client system and the corresponding one of the first limit value and second limit value.</p>	See Claim 2.
<b>Claim 4</b>	
<p>4. The method of claim 2 wherein step (iv) comprises the steps of:</p>	See Claim 2.
<p>(a) the user selecting one of the first price and second price of the current</p>	See Claim 2.

Exhibit E

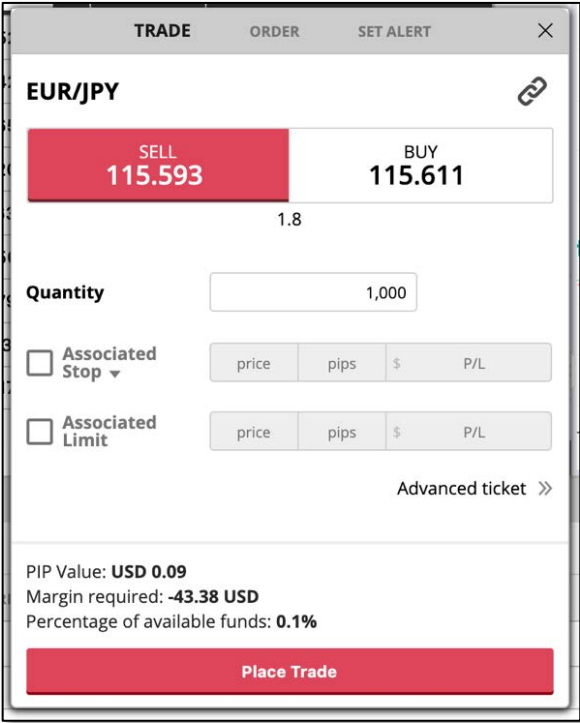
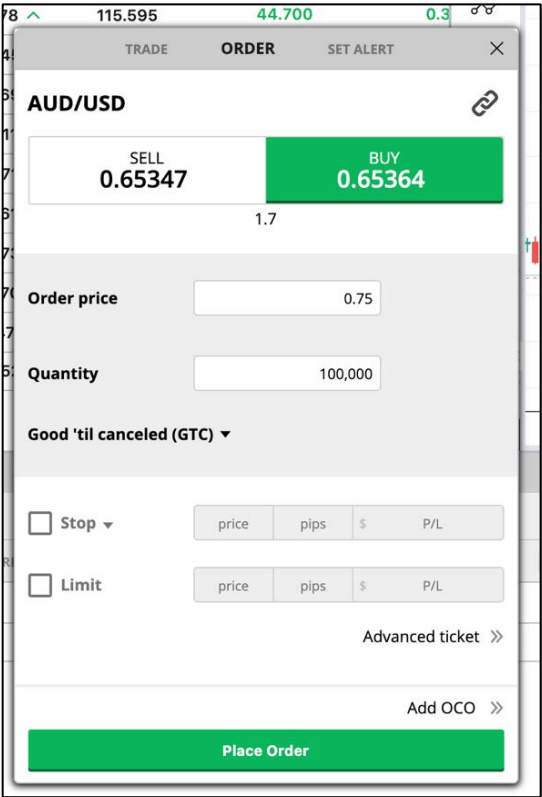
<p>exchange rate displayed at the trading client system;</p>	
<p>(b) displaying to the user a set of input fields to define a desired trade, the input fields including an identification of the pair of currencies the user desires to trade, the amount of the currencies desired to be traded, the selected first price or second price of the current exchange rate received at the trading client system and a limit value, and where the input fields to identify the pair of currencies and the first price or second price are populated with appropriate values determined from the user's selection of the one of the first price or second price;</p>	<p>See:</p> 

Exhibit E



(c) receiving from the user input to the input field defining the desired amount of currency to be traded; and

See above. See also:

Exhibit E

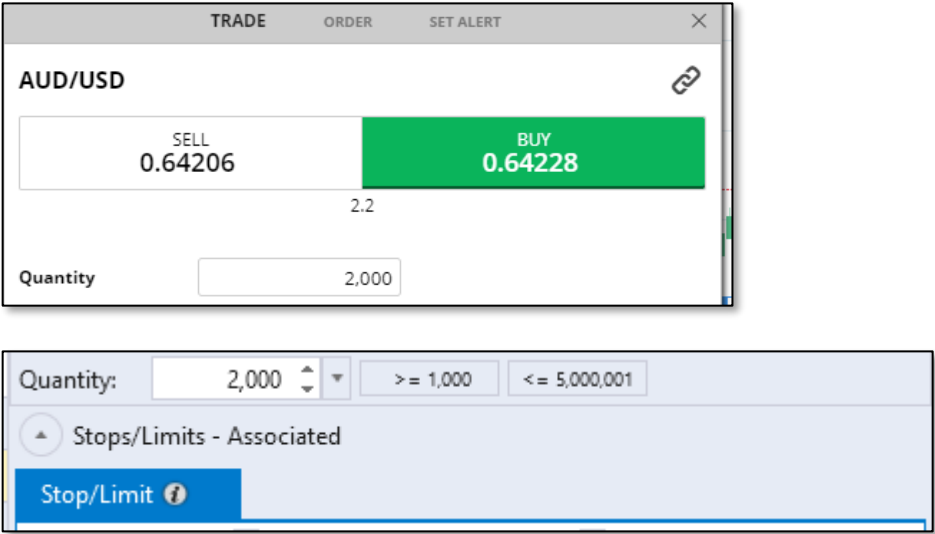
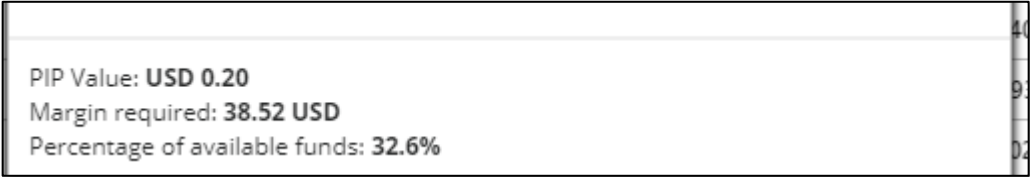
	
<p>(d) determining the requested trade price from the selected one of the first price and second price and the limit value.</p>	<p>See above. See also:</p> 
<p><b>Claim 5</b></p>	
<p>5. The method of claim 4 wherein in step (b) the displayed set of input fields includes: a first limit value to define a maximum acceptable difference between the first price of the current exchange rate received at the trading client system and the first price of the</p>	<p>See above. See also:</p>

Exhibit E

corresponding current exchange rate determined at the trading client system; and a second limit value to define a maximum acceptable difference between the second price of the current exchange rate received at the trading system and the second price of the corresponding current exchange rate determined at the trading client system and in step (d) the requested trade price is derived from the selected first price or second price and the corresponding one of the first limit value and second limit value.

The screenshot shows a trading interface for EUR/JPY. At the top, there are tabs for 'TRADE', 'ORDER', and 'SET ALERT'. Below this, there are four order buttons: 'SELL 115.602' (white), 'BUY 115.618' (green), 'SELL 115.602' (red), and 'BUY 115.618' (white). A '1.6' pip value is indicated between the first and second buttons, and between the third and fourth buttons. A note states: 'These orders are linked. When one is executed the other will be cancelled.' Below the buttons are two identical order entry forms. Each form has fields for 'Order price', 'Quantity' (set to 1,000), and 'Good 'til canceled (GTC)'. There are also checkboxes for 'Stop' and 'Limit' orders, each with a corresponding input field for quantity, price, pips, \$, P/L, and GTC. At the bottom of the interface, there is a 'Standard ticket' link and a 'Remove OCO' button. A 'Enter a price' input field is also visible.

PIP Value: **USD 0.20**  
 Margin required: **38.52 USD**  
 Percentage of available funds: **32.6%**

Exhibit E

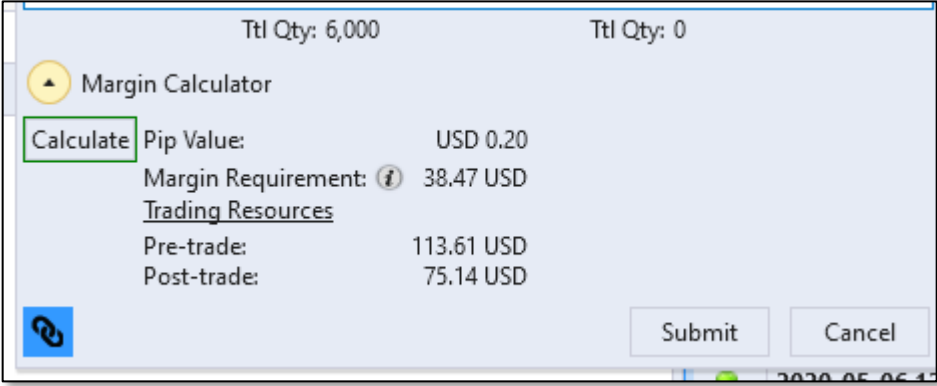
	
<p><b>Claim 6</b></p>	
<p>6. The method of claim 2 wherein, when the limit value is zero, the requested trade price is the current corresponding first price or second price of the current exchange rate at the trading server.</p>	<p><i>See generally, Claim 1.</i></p>
<p><b>Claim 7</b></p>	
<p>7. A method of trading currencies over a computer network connecting a trading system server and at least one trading client system, comprising the steps of:</p>	<p><i>See Claim 1.</i></p>
<p>(i) at the trading system server, determining and dynamically maintaining a plurality of current exchange rates, each current exchange rate relating to a</p>	<p><i>See Claim 1(i).</i></p>

Exhibit E

<p>pair of currencies and including a first price to buy a first currency of the pair with respect to a second currency of the pair and a second price to sell the first currency of the pair with respect to the second currency of the pair;</p>	
<p>(ii) transmitting data from the trading system server to a trading client system, the transmitted data representing at least one current exchange rate at the time of the transmission;</p>	<p>See Claim 1(ii).</p>
<p>(iii) receiving at the trading system server input from a user of the trading client system identifying a pair of currencies the user desires to trade, an amount of at least one currency of the pair desired to be traded and a requested trade price at which it is desired to effect the trade;</p>	<p>See Claim 1(iii-iv). <i>See also:</i></p>

Exhibit E

The screenshot displays a trading platform interface. On the left, an 'Orders (3)' window shows a table of active orders:

MARKET	ORDER	
EUR/USD	Sell	2,000 Delete
AUD/USD	Buy	2,000 Delete
AUD/USD	Buy	2,000 Delete

Below the table is a candlestick chart for EUR/USD from December 2020 to April 2021. A red dashed horizontal line is drawn at a price level of 1.08150. An 'EDIT ORDER' window is overlaid on the right, showing the following details:


- Pair: EUR/USD
- Direction: SELL (highlighted in red)
- Price: 1.08150
- Quantity: 2,000
- Order Type: Good 'til canceled (GTC)
- Stop and Limit options are visible but not selected.
- Buttons: Delete Order, Update Order

Active Orders													
Market	Market Expiry	Buy/Sell	Order Type	Quantity	Order Price	Current Price	Distance Away	Stop	Limit	Expiry	Order ID	Date/Time	Basis
EUR/USD		Sell	Limit	2,000	1.1	1.08143	0.01857			Good Till Cancelled	732593048	2020-05-06 11:20:00 AM	Single
AUD/USD		Buy	Limit	2,000	0.64	0.64198	0.00198			Good Till Cancelled	732606240	2020-05-06 12:04:41 PM	Single
AUD/USD		Buy	Limit	2,000	0.64	0.64198	0.00198			Good Till Cancelled	732606441	2020-05-06 12:05:22 PM	Single
AUD/USD		Sell	Associated Stop	4,000	0.641	0.64177	0.00077			Good Till Cancelled	732608284	2020-05-06 12:13:09 PM	Single

*Oanda Corp. v. GAIN Capital Holdings, Inc.;*  
*GAIN Capital Group, LLC.*



Exhibit E

<p>(iv) at the trading system server, comparing the requested trade price to the respective first price or second price of the corresponding current exchange rate at that time and, if the respective first price or second price of the corresponding current exchange rate at that time is equal to or better than the requested trade price, effecting the trade at the corresponding respective current exchange rate first price or second price and if the corresponding current exchange rate is worse than the requested trade price, refusing the trade; and</p>	<p>See above. See also Claim 1(iv-v).</p>																		
<p>(v) transmitting from the trading system server to the trading client system an indication of whether the trade was refused or transacted and, if transacted, an indication of the price the trade was transacted at.</p>	<p>See above. See also Claim 1(v-vii).</p>  <p>The screenshot shows a news feed with the following entries:</p> <table border="1"> <thead> <tr> <th>Time</th> <th>News Title</th> </tr> </thead> <tbody> <tr> <td>2020-05-06 12:02:30 PM</td> <td>Italy's daily coronavirus death toll and new cases' tally both rise</td> </tr> <tr> <td>2020-05-06 12:02:23 PM</td> <td>*TOP NEWS*-Central Banks &amp; Global Economy</td> </tr> <tr> <td>2020-05-06 12:02:21 PM</td> <td>*TOP NEWS*-Global Markets</td> </tr> <tr> <td>2020-05-06 12:00:32 PM</td> <td>*TOP NEWS*-World News</td> </tr> <tr> <td>2020-05-06 11:58:18 AM</td> <td>REFILE-UPDATE 2-Sterling hits 12-day low versus dollar</td> </tr> <tr> <td>2020-05-06 11:57:47 AM</td> <td>UPDATE 1-Natixis swings to a loss, hikes provisions on coronavirus, frauds</td> </tr> <tr> <td>2020-05-06 11:55:51 AM</td> <td>BUZZ-COMMENT-Yen gains counter upbeat risk, AUD/USD rally struggles</td> </tr> <tr> <td>2020-05-06 11:55:43 AM</td> <td>WRAPUP 2-U.S. private navrolls dive by a record 20.2 million</td> </tr> </tbody> </table> <p>At the bottom of the screenshot, there is a trade confirmation pop-up:</p> <p>New trade: Market AUD/USD, BUY, 0.64149, Quantity 2,000, Order ID [REDACTED]</p>	Time	News Title	2020-05-06 12:02:30 PM	Italy's daily coronavirus death toll and new cases' tally both rise	2020-05-06 12:02:23 PM	*TOP NEWS*-Central Banks & Global Economy	2020-05-06 12:02:21 PM	*TOP NEWS*-Global Markets	2020-05-06 12:00:32 PM	*TOP NEWS*-World News	2020-05-06 11:58:18 AM	REFILE-UPDATE 2-Sterling hits 12-day low versus dollar	2020-05-06 11:57:47 AM	UPDATE 1-Natixis swings to a loss, hikes provisions on coronavirus, frauds	2020-05-06 11:55:51 AM	BUZZ-COMMENT-Yen gains counter upbeat risk, AUD/USD rally struggles	2020-05-06 11:55:43 AM	WRAPUP 2-U.S. private navrolls dive by a record 20.2 million
Time	News Title																		
2020-05-06 12:02:30 PM	Italy's daily coronavirus death toll and new cases' tally both rise																		
2020-05-06 12:02:23 PM	*TOP NEWS*-Central Banks & Global Economy																		
2020-05-06 12:02:21 PM	*TOP NEWS*-Global Markets																		
2020-05-06 12:00:32 PM	*TOP NEWS*-World News																		
2020-05-06 11:58:18 AM	REFILE-UPDATE 2-Sterling hits 12-day low versus dollar																		
2020-05-06 11:57:47 AM	UPDATE 1-Natixis swings to a loss, hikes provisions on coronavirus, frauds																		
2020-05-06 11:55:51 AM	BUZZ-COMMENT-Yen gains counter upbeat risk, AUD/USD rally struggles																		
2020-05-06 11:55:43 AM	WRAPUP 2-U.S. private navrolls dive by a record 20.2 million																		