Dynamic Θ Time:

Algorithm, Configuration, Tests

Wolfgang Breymann

WAB.2000-07-31

January 6, 2009

An internal paper by the Olsen Ltd. Research Group

Abstract

The Dynamic Θ Time is a tool for deseasonalizing intraday financial time series. Its main advantages, compared to the static ϑ time developed at O&AS are universability, configurability, and adaptability. The Dynamic Θ Time works for delocalized over-the-counter (OTC) markets as well as for localized exchanges with well-defined opening and closing hours. In particular, it can deal with very abrupt changes in the activity pattern, it takes into account regional market components which can be freely configured, and it adapts to changes in the activity pattern.

The algorithm is based on market activity histogram computed by intra-week moving averages. This assures good modelization properties and adaptability. In addition, regional public holidays and the effect of Daylight Saving Time on the activity pattern is taken into account.

Tests of the deseasonalization quality show only very few remaining seasonal components in the volatility.

Keywords: Financial time series, volatility, seasonalities, intra-day data, time transformation, public holidays, Daylight Saving Time.

Contents

List of Figures

List of Tables

1 Introduction:
A General Adaptive Volatility-Driven Time Scale

This document describes the algorithm of the *Dynamic* Θ *Time*, together with two examples and results of the deseasonalization quality. The incentive for developing the *Dynamic* Θ *Time* and its rationale have been discussed in detail in (Breymann, 1998). A summary of the algorithm and some results have been presented in (Breymann *et al.*, 2000). The algorithm is flexible, adaptive, and universally applicable to arbitrary high-frequency financial instruments. It uses and extends the newest Olsen & Associates basis technology as Exponential Moving Averages (EMA) and iterated EMA (Zumbach and Müller, 2001) and takes into account the internal structure of financial markets in various ways.

To fix the ideas, the general problem is briefly described. A given (continuous or discrete) financial time series

$$x_t \tag{1}$$

will typically exhibit deterministic seasonalities together with other (deterministic or stochastic) features. The goal is to separate the seasonalities from the other features, i.e., to *deseasonalize* the time series. This deseasonalization is a prerequisite if one wants to extract more subtle effects from (1).

In financial time series the most prominent seasonal pattern is the volatility. As example the weekly volatility pattern for EUR/USD spot rates is shown in Fig. 1. A varieties of measures are in use for getting quantitative volatility estimates. Here, we define the volatility $v[\Delta T](t)$ of a time series $x(t)$ with respect to a time horizon ΔT as a kind of 1-norm:

$$v[\Delta T](t) = \langle |x(t) - x(t - \Delta t)| \rangle, \tag{2}$$

where $\langle\rangle$ denotes a suitable way of averaging explained below (cf. section 2.2).

The time series $x(t)$ depends on the instrument under consideration. For the FX spot market, e.g., it is the logarithmic middle price

$$x(t) = \frac{\log p_{\text{bid}}(t) + \log p_{\text{ask}}(t)}{2} \tag{3}$$

and for a stock index, it is the logarithmic index.

Even though from a conceptual point of view, there is no difference between deseasonalizing the volatility or any other time-dependent quantity related to the time series (1), we will focus here on the volatility. However, it is pointed out that the Dynamic Θ Time is designed in a way that it can be used for other quantities with minor changes only.

Fig. 2 displays the daily volatility pattern for a number of instruments. The patterns exhibit a variety of different shapes, depending on the nature of the financial asset (FX, bonds, equities, commodities), the kind of instruments (spot, future, etc.) and the organization of the market (OTC vs. localized, exchange-traded markets, open outcry vs. computerized trading). The extreme patterns are

(i) a rather smooth pattern with broad maxima, which is typical for delocalized OTC markets with round-the-clock activity as in the case of major FX spot rates, and

(ii) the well-known U-shape pattern typical for localized, exchange-traded markets, which is mainly due to the well-defined opening and closing time.

Between these two extreme cases there exists a variety of transition patterns. The Dynamic Θ Time is designed in a way that it can handle all these different instruments.

A time transformation similar to the Dynamic Θ Time has previously been developed by O&A for the FX spot market. This algorithm, which is described in (Müller *et al.*, 1990), will be refered to as static ϑ time. The authors first proposed the relation between time transformation, market activity, and volatility exposed in Sec. 2.1. Then, they used a superposition of 7th and 5th-order polynomials to fit the market activity of three regional components to the observed volatility pattern. Because of this modelling decision, the static ϑ time is not adaptive, and it is not very suitable for localized, exchange-traded instruments. The deseasonalization quality achieved with the Dynamic Θ Time is better than that achieved with static ϑ time even for FX spot rates (cf. Sec. 5).

This paper is organized as follows. Section 2 gives a detailed account of the algorithm of the Dynamic Θ Time. The core of the algorithm is the computation of the weekly market activity pattern (Sec. 2.2). The relation between time scale, market activity (Sec. 2.1), and volatility (cf. Sec. 2.2.1) is

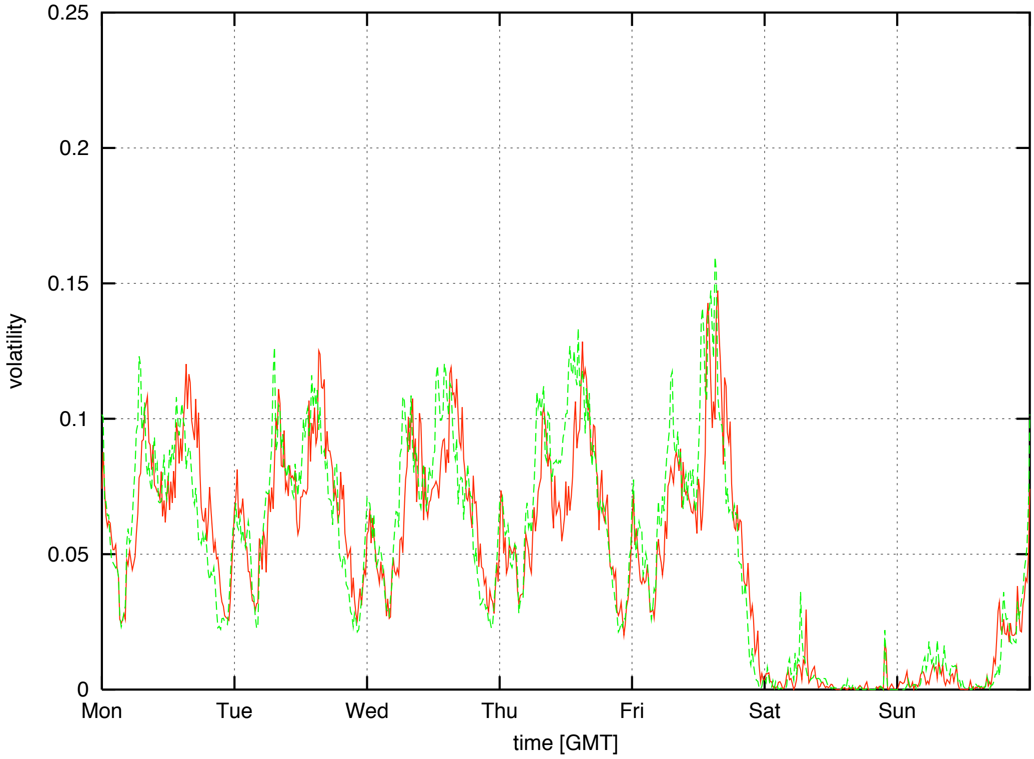

Figure 1: Weekly volatility pattern for the EUR/USD spot rate, averaged over the years 1997/1998. The time on the abscissa is given in hours GMT. Full, red line: summer time. Dashed, green line: winter time.

presented, the periodic conditional average (PCA), the intra-week moving average (IWMA) and the computation of the activity histograms is exposed in detail (Sections 2.2.2 and 2.2.3), and the update mechnism is explained (Sec. 2.2.4). The exposition of the activity patterns concludes with remarks on Daylight Saving Time (Sec. 2.2.5). In Sec. 2.3 it is explained how the market is decomposed into regional components and how by means of this decomposition public holidays and so-called partial and fuzzy holidays can be taken into account very accurately. Section 3 lists the most important differences between the Dynamic Θ Time and the static ϑ time. In Sec. 4, the main pecularities of the Dynamic Θ Time are explained for representatives of two important groups of instruments: major FX spot rates as examples for delocalized, OTC instruments, and the German DAX as an example of a localized, exchange traded instrument. For these instruments, tests of the deseasonalization quality are presented in Sec. 5.

In addition, the most important properties of exponential moving averages (EMA) and their iterates are summarized in App. A, a summary of the Dynamic Θ Time algorithm is presented in App. B, and the configuration parameters of the Dynamic Θ Time are explained in App. C.

2 The Algorithm

2.1 Dynamic Θ Time and Market Activity

The basic task of the Dynamic Θ Time is to transform physical time to an appropriate time scale driven by the volatility observed for a given security. The starting point is the definition of the market-activity driven time scale defined by (Dacorogna *et al.*, 1993)

$$\Theta(t_2) - \Theta(t_1) = \Delta\Theta(t_1, t_2) = \int_{t_1}^{t_2} a(t)dt, \tag{1}$$

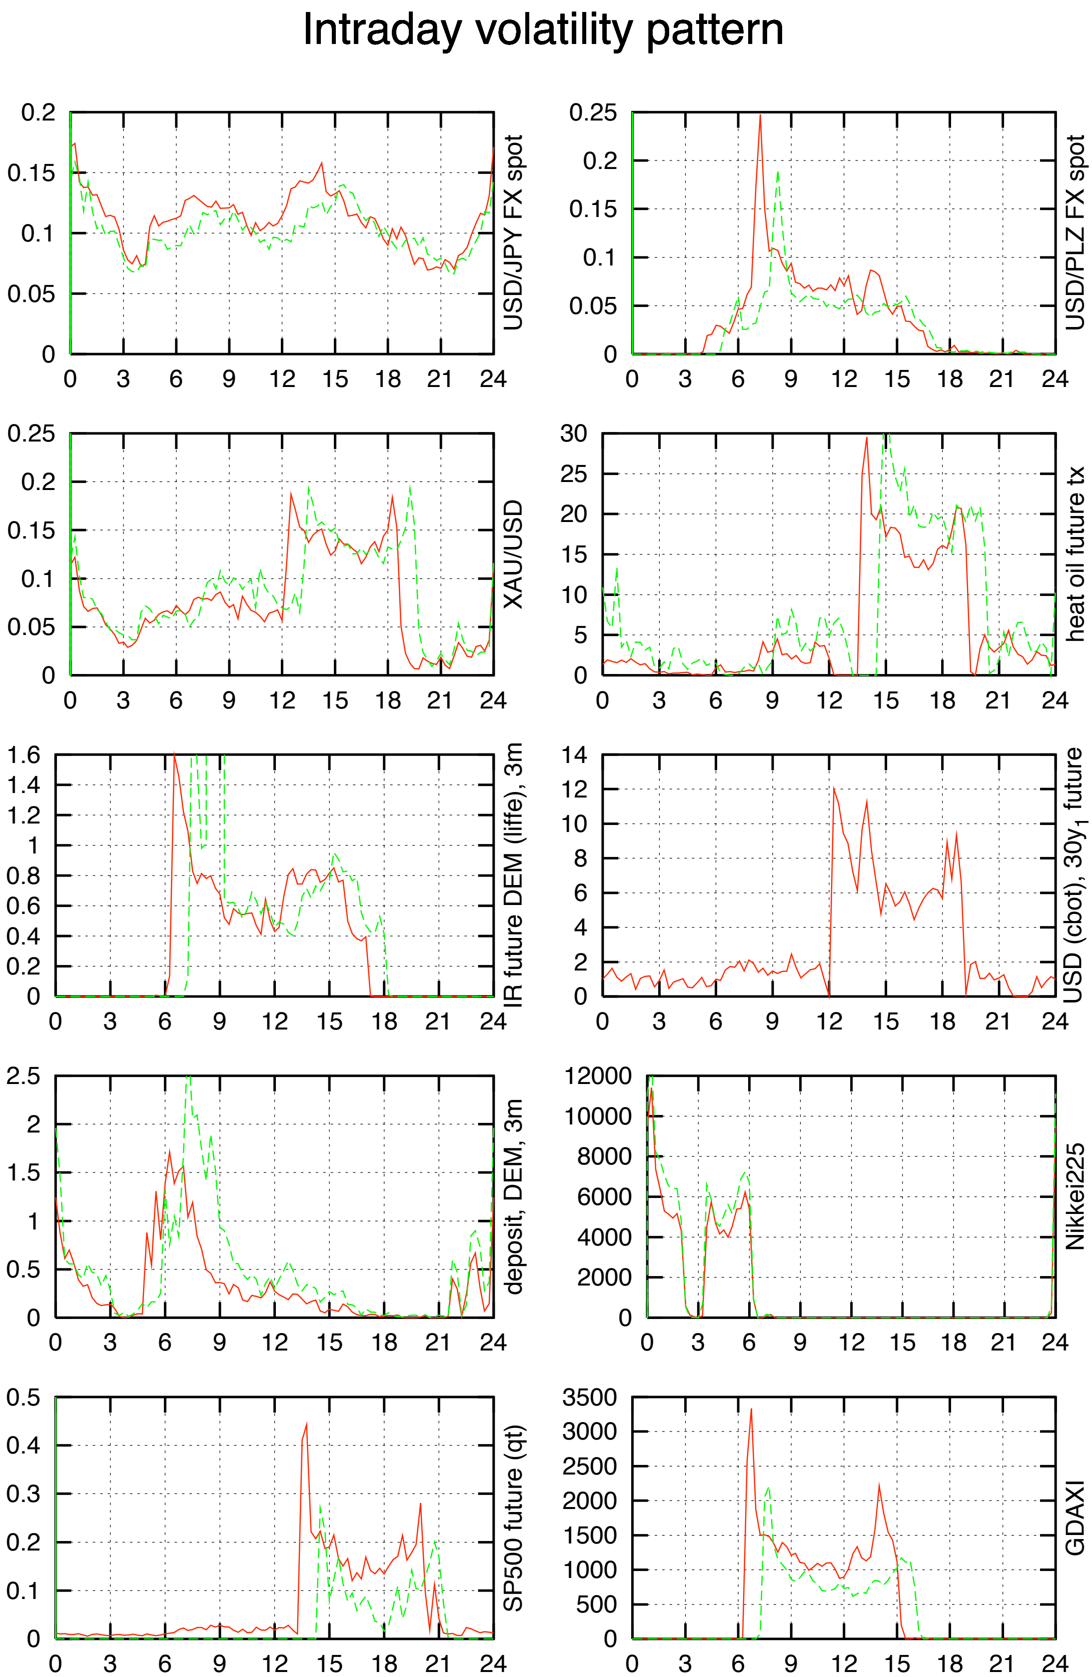

Figure 2: Daily volatility pattern for some selected instruments, averaged over the years 1997/1998. The time on the abscissa is given in hours GMT. Full, red line: summer time. Dashed, green line: winter time. Notice that the averaging period starts after the extension of the floor trading hours for the German DAX (lower right).

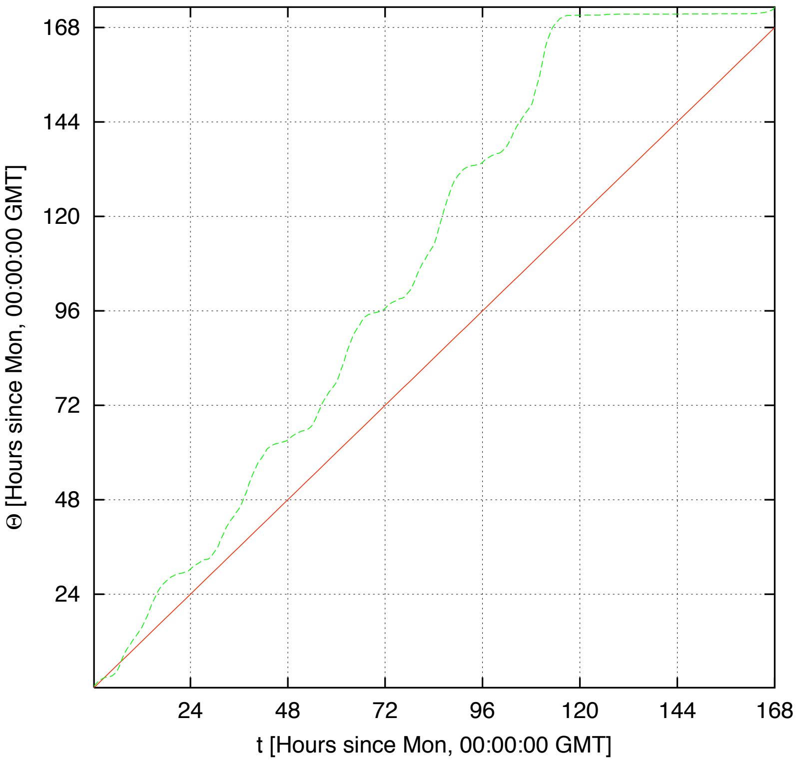

Figure 3: Dynamic Θ Time as function of physical time for EUR/USD for a week without public holidays, start Monday, 00:00:00 GMT. 1 week in physical time corresponds to about 3% more than 1 week in Dynamic Θ Time in order to compensate for the reduction in market activity during public holidays.

where the *market activity* $a(t) \geq 0$ is a quantity indicating the *conditional intra-week average* flow of *relevant events* in a financial market, which define the flow of Θ-Time through physical time, $\mathrm{d}\Theta/\mathrm{d}t$ (cf. Fig. 3). The meaning of 'conditional intra-week average' will be defined below.

Good modeling of the market activity is essential for the quality of the Dynamic Θ Time scale. The important issues are:

- The choice of an appropriate quantity serving as proxy for market activity. Unfortunately, the term 'relevant events' is somewhat unprecise but there are a number of quantites which may fulfil our requirements, as the tick rate or the volatility. For stability reasons we have chosen volatility (cf. next paragraph), but extensions to other quantities are straight forward.
- The possibility to model arbitrary activity patterns with abrupt changes in market activity. This is important because in local, exchange-traded markets with well-determined opening and closing time the activity pattern suddenly jumps from zero to a positive value at market opening and it jumps back to zero at market closing. Notice that such patterns cannot adequately be modeled by 7th order polynomials, as used in (Dacorogna *et al.*, 1993).
- The modeling of the slow dynamic behavior of the activity pattern. This is important because market activity slowly changes with time, due to changes in the habits of the market participants, institutional changes, etc. This is why a special kind of a moving average, the *intra-week moving average* (IWMA) is used (cf. section 2.2.2).
- The modeling of the geogaphical composition of market activity. This makes it possible to take into account perturbations of the regular weekly activity pattern due to local holidays affecting only parts of the market.
- Keeping track of Daylight Saving Time. This is necessary because the average activity pattern suddenly changes during the transition between winter time and summer time. Different DST

periods are accounted for by treating their activity pattern separately.

For 'normal' days, i.e., when there is no public holiday in any market component, the activity $a(t)$ is set to the conditional intra-week averaged activity computed by means of the IWMA plus a small constant positive component:

$$a(t) = \mathrm{IWMA}[\tau; a_\alpha](t) + a_0. \tag{2}$$

The symbols have the following meaning:

- α labels the different DST periods,
- $\mathrm{IWMA}[\tau; \cdot]$ denotes the intraweek moving average with range τ,
- $a_\alpha(t)$ is the full market activity pattern with respect to the Daylight Saving Period α,and
- a_0 is a small positive constant assuring that the market activity is always strictly positive.[1]

The intra-week moving average is performed with respect to the Daylight Saving period containing t.

When there is a holiday in any of the market components, the activity of this component must be discounted. This is the case for roughly 20 days per year. Discounting is achieved by modifying eq. (2) in the following way. The market is decomposed into regional components $i = 1 \ldots n$ such that the total activity $a(t)$ is given as the sum of component activities $a_i(t)$,

$$a(t) = a_0 + \sum_{i=1}^{n} a_i(t), \tag{3}$$

The number of components as well as their geographical location can be freely configured by the user. This is important mainly for instruments traded in localized exchanges as equities, futures, and standard option. But also for some OTC instrument a modification of the standard components turns out to be useful (cf. Sec. 4).

The single-component activity is modelled as

$$a_i(t) = \Big(\frac{s_0}{n} + h_i(t)\, s_i(t)\Big)\, \mathrm{IWMA}[\tau; a_\alpha](t), \tag{4}$$

where

- $s_i(t)$, $i = 1 \ldots n$ determines the momentary share of the ith market component, s_0 being a constant factor, and
- $h_i(t)$, $i = 1 \ldots n$ is the holiday factor of market component i, which assumes values between 0 and 1.

The constant share s_0 accounts for market components not modelled explicitly. The components' activity shares observe the normalization condition

$$\sum_{i=0}^{n} s_i = 1, \tag{5}$$

and $h_i(t)$ is 1 if in component i the corresponding day is a normal day (working day or weekend) and 0 if it is a holiday.

Inserting eq. (4) into (3) yields for the total activity:

$$a(t) = a_0 + \left(s_0 + \sum_{i=1}^{n} h_i(t)\, s_i(t)\right) \mathrm{IWMA}[\tau; a_\alpha](t). \tag{6}$$

It can easily be checked that eq. (6) do indeed reduce to eq. (2) if $h_i(t) = 1$ for all $i = 1 \ldots n$.

The modelling of the activity $\mathrm{IWMA}[\tau; a_\alpha](t)$ is described in the next section, and the activity shares $s_i(t)$ and the holiday factors $h_i(t)$ are described in Sections 2.3.1 and 2.3.2, respectively.

2.2 The Weekly Activity Pattern

The important points are the relation between market activity, volatility, the intra-week moving average, and the update mechanism.

[1]Strict positivity of the market activity is a technical condition necessary for the inversion of the Dynamic Θ Time transformation.

2.2.1 Market Activity and Volatility

The mean weekly activity[2] $\mathrm{IWMA}[\tau; a_\alpha](t)$ is calculated from the mean volatility $\mathrm{IWMA}[\tau; v_\alpha](t)$ pattern aggregated over a period of several month (cf. Sec. 2.2.4).

In (Dacorogna *et al.*, 1993) it had been argued that the activity is related to the volatility (2) through a scaling law,

$$a[\Delta T] \sim (v[\Delta T])^\gamma \tag{7}$$

with scaling exponent $\gamma = E$, where $1/E$ is the drift exponent governing the scaling law of the volatility as a function of the time horizon ΔT (Müller *et al.*, 1990):

$$v[\Delta T] \sim (\Delta T)^{1/E}. \tag{8}$$

While for major FX rates, (Müller *et al.*, 1990) report values of $1/E \simeq 0.58$ corresponding to $\gamma \simeq 1.72$, in the case of normal diffusion $\gamma = 2$ is expected. Best deseasonalization results have indeed been obtained for $\gamma \simeq 2$ so that this value is used instead of E.[3] Hence

$$\mathrm{IWMA}[\tau; a](t) = c \cdot \mathrm{IWMA}[\tau; v[\Delta T]]^2(t) \tag{9}$$

The normalization factor is chosen such that in average, Dynamic Θ Time passes as fast as physical time. In practice it is determiend by the condition that the interval in Dynamic Θ Time corresponding to a 4-year reference period $[T_0, T_1]$ should be of exactly the same length:

$$T_1 - T_0 \stackrel{!}{=} \int_{T_0}^{T_1} a(t)\,\mathrm{d}t.$$

2.2.2 Periodic Conditional Averages

The activity pattern is computed by a *periodic conditional average* (PCA), which makes it possible to extract a seasonal pattern of period T_p from a noisy time series x_t. In general the periodic conditional average is defined as the expectation conditional to the time $t_p = t \bmod T_p$ within a period of length T_p:

$$\mathrm{PCA}[T_p; x](t_p) = E[x_t | t \text{ mod } T_p = t_p] \tag{10}$$

For a PCA over a period $[T_{\min}, T_{\max}]$ this can be rewritten as

$$\begin{aligned} \mathrm{PCA}[T_p; x] &= C \int_{T_{\min}}^{T_{\max}} x(t') \sum_{n=-\infty}^{\infty} \delta(t' - nT_p)\,\mathrm{d}t' \\ &= C \sum_{nT_p \in [T_{\min}, T_{\max}]} x(nT_p) \end{aligned} \tag{11}$$

where $\delta(t)$ denotes Dirac's δ function and the normalization constant

$$C^{-1} = \int_{T_{\min}}^{T_{\max}} \sum_{n=-\infty}^{\infty} \delta(t' - nT_p)\,\mathrm{d}t' \tag{12}$$

counts the number of peaks in interval $[T_{\min}, T_{\max}]$. Notice that due to the sum of δ-functions the integrals in eqs. (11) and (12) reduce to a finite sum.

Eq. (11) can be generalized to allow for the use of a kernel $\omega(t)$:

$$\begin{aligned} \mathrm{PCA}[T_p; x](t) &= C \int_{-\infty}^{\infty} x(t')\,\omega(t - t') \sum_{n=-\infty}^{\infty} \delta(t' - nT_p)\,\mathrm{d}t' \\ &= C \sum_{n=-\infty}^{\infty} x(t - nT_p)\,\omega(nT_p), \end{aligned} \tag{13}$$

$$C^{-1} = \sum_{n=-\infty}^{\infty} \omega(nT_p). \tag{14}$$

To assure adaptability, the PCA is combined with the exponential moving averages (EMA) or moving averages (MA) exposed in (Zumbach and Müller, 2001) by choosing a corresponding kernel in

[2] Even though, strictly speaking, this quantity is not an IWMA, this notation is used for simplicity.

[3] The Dynamic Θ Time allows the user to choose other definitions for the volatility, in which case the values of γ must be adjusted accordingly.

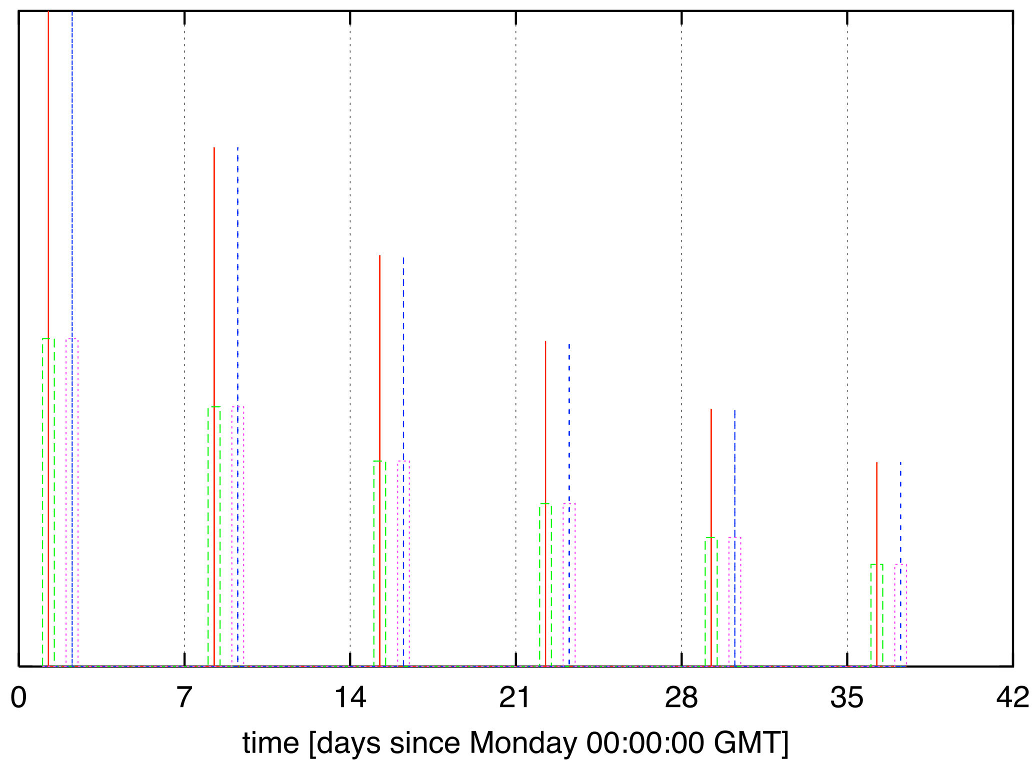

Figure 4: Illustration of the kernel for the PCEMA (16) with $k = 1$ and for the effective PCEMA kernel used for the computation of the volatility histograms. Periodicity is $T_p = T_w \equiv 1\,\text{week}$. For a good visibility the bin width is set to 12 hours. Red and blue lines: PCEMA kernel (15) for $t \mod 1\text{week} = 30$ hours and 54 hours, respectively. Green and pink lines: Effective kernels entering the evaluation of the volatility histograms, as exposed in Sec. 2.2.3.

eq. (13). Accordingly, an *periodic conditional iterated exponential moving average* (PCEMA) of the time series $x(t)$ is defined as

$$\begin{aligned}\text{PCEMA}[\tau, T_p, k; x](t) &= C_k \int_{-\infty}^{t} x(t') \left(\frac{t-t'}{\tau}\right)^{k-1} e^{-\frac{t-t'}{\tau}} \sum_{n=-\infty}^{\infty} \delta(t - t' - nT_p)\, \mathrm{d}t' \\ &= \sum_{n=0}^{\infty} x(t - nT_p)\, \text{pcema}[\tau, T_p, k]\,(n) \qquad (15)\end{aligned}$$

with the periodoc conditional iterated EMA kernel

$$\text{pcema}[\tau, T_p, k]\,(n) = C_k \left(\frac{nT_p}{\tau}\right)^{k-1} e^{-\frac{nT_p}{\tau}}. \qquad (16)$$

and the normalization factor C_k being determined by the relation

$$C_k^{-1} = \left(\frac{T_p}{\tau}\right)^{k-1} \sum_{n=0}^{\infty} n^{k-1} x(t - nT_p) e^{-\frac{T_p}{\tau} n}. \qquad (17)$$

Notice that because of the intra-week conditional average, the normalization constant $C \neq 1$ is different from that of the iterated EMA defined in eq. (A.6) of the appendix.

In analogy to eq. (A.11) of App. A the mth order approximation of a *periodic conditional moving average* (PCMA) is defined as

$$\text{PCMA}[\tau, T_p, m; x](t) = \sum_{n=0}^{\infty} x(t - nT_p)\, \text{pcma}[\tau, T_p, m]\,(n) \qquad (18)$$

with

$$\mathrm{pcma}[\tau, T_p, m]\,(n) = \frac{1}{m} \sum_{k=1}^{m} \mathrm{pcema}[\tau', T_p, k]\,(n) \tag{19}$$

and $\tau' = 2\tau/(m+1)$, as for the moving average operator defined by (Zumbach and Müller, 2001).

In the following, only the case of $T_p = T_w \equiv 1\,\text{week}$ is considered, and the various PCAs are called *intra-week averages*. Thus, Eq. (18) in connection with eq. (19) and the identity $\mathrm{iwma}[\tau, m] = \mathrm{pcma}[\tau, T_w, m]$ are the defining equations of the IWMA of order m. The definition contains a discrete sum instead of an integral, which is an immediate consequence of the conditioning operator. Therefore, IWMAs can be computed for regular discrete time series without approximation if the time step of the series is commensurable with 1 week.

2.2.3 Volatility and Activity Histograms

For practical computation of intra-week averages, a weekly pattern is evaluated only for a discrete number of points. The computation can be conceived as a three-step procedure, which is described here for the case of the volatility. In the first step a regular discrete volatility series $v_{\mathrm{reg}}[\Delta T](t_i)$ with time step δ is computed from the incoming price series at time points $t_i = i\delta$, with $i \in N$ and time horizon $\Delta T = 2\delta$. In the second step, this time series is slightly smoothed by means of the following kernel:

$$v_{\mathrm{sm}}(t_i) = \frac{1}{\delta} \int_{t_i-\delta}^{t_i+\delta} \omega(t' - t_i)\; \mathrm{MA}[\tau_s, v_{\mathrm{reg}}[\Delta T]](t')\,\mathrm{d}t' \tag{20}$$

with a MA of a short range τ_s and the tent-shaped kernel of width 2δ defined as

$$\omega(t) = \begin{cases} 1 + \dfrac{t}{\delta} & \text{if } -\delta < t \le 0 \\[2ex] 1 - \dfrac{t}{\delta} & \text{if } 0 < t \le \delta \\[2ex] 0 & \text{elsewhere.} \end{cases} \tag{21}$$

In the third step, the IWMA is computed separately for every bin. According to eq. (18) this yields

$$\mathrm{IWMA}[\tau, m; v_{\mathrm{sm}}](t_i) \equiv \mathrm{IWMA}[\tau, m; v_{\mathrm{sm}}](t_i) = \sum_{n=0}^{\infty} v_{\mathrm{sm}}(t_i - n)\; \mathrm{iwma}[\tau, m]\,(n). \tag{22}$$

It is convienent to store the values of the last period in a histogram of total width T_w and with bins of width Δ. In the present case the convention is adopted that the pattern should start at Monday, 00:00:00 GMT. Then the time axis of the histogram represents the time in the week, $t_w = t \bmod 1\,\text{week}$, and the bins are centered at

$$t_{w,i} = \left(i - \frac{1}{2}\right)\Delta, \qquad i = 1, \ldots, \frac{T_w}{\Delta}.$$

Finally, the activity histogram is computed from the volatility pattern by means of eq. (9).

Numerical values of δ, τ_s, and Δ are listed for some examples in Tab. 3.

2.2.4 The Update Mechanism

Volatility histogram. A tick by tick update of the volatility histogram is possible thanks to the iterative formula (A.3) of App. A for the intra-week iterated EMA. In fact, the continuous update is a three-step procedure according to the three-step procedure for the computation of the histograms. In the first step the incoming tick is used to compute the next value of the regularized volatility series $v_{\mathrm{reg}}(t_i)$. It is at this stage that interpolation of incoming data is performed if needed. There are two possible interpolation scheme: Linear interpolation and interpolation according to the diffusion-like scaling law, i.e., such that the interpolated absolute returns scale with exponent 0.5 as a function of time.

In the second step, this series is slightly smoothed by means of eq. (20). In the third step, the histogram value is updated by

$$\begin{aligned} \mathrm{IWEMA}[\tau', 1; v_{\mathrm{sm}}](t_i) &= \mu\, \mathrm{IWEMA}[\tau', 1; v_{\mathrm{sm}}]\,(t_i - T_w) + (1-\mu)\, v_{\mathrm{sm}}(t_i) \\ \mathrm{IWEMA}[\tau', k; v_{\mathrm{sm}}](t_i) &= \mu\, \mathrm{IWEMA}[\tau', k; v_{\mathrm{sm}}](t_i - T_w) + (1-\mu)\, \mathrm{IWEMA}[\tau', k-1; v_{\mathrm{sm}}](t_i), \\ & \qquad k = 2, \ldots, m \end{aligned} \tag{23}$$

where the $\mathrm{IWEMA}[\tau', k; v_{\mathrm{sm}}]$, $k = 1, \ldots m$ are the IWEMAs composing the IWMA according to eq. (19) with τ' as defined below that equation and t_i is the time in the week coresponding to bin i. A range of $\tau_{\mathrm{IWMA}} = 1\,\mathrm{month}$ is used for the IWMA (cf. section C). Since interpolation has already been taken into account in the first step, no interpolation scheme is needed at this point.

Activity histogram. Activity histograms have the same binning as volatility histograms. They are calculated by means of eq. (9) for any bin. To respect the periodicity of the weekly pattern, the activity histogram is updated from the volatility histogram always Mondays , 00:00:00 GMT.
Remark. The weekly histograms automatically account for day-of-the-week effects. The computation of the weekly seasonality pattern is both fully adaptive and very flexible and does not depend on instrument-specific parameters.

2.2.5 Daylight Saving Time

Daylight Saving Time is taken into account at the level of the market activity histograms. The main effect of Daylight Saving Time is a one-hour shift of the activity and volatility patterns for the affected components. Since there is no Daylight Saving Time in East-Asian markets, the shift affect only parts of the pattern (cf. left column, patterns 1 and 2 from top in Fig. 2). Since these partial shifts are difficult to disentangle, the most straightforward way to properly deal with them is to compute the activity and volatility histogram separately for DST and non-DST periods.

In practice, only the activity histograms are stored separately for the DST and the non-DST period. At the beginning of any new such period, the volatility histogram is initialized from the corresponding activity histogram by mean of the inverse of eq. (9). In the sequel, the volatility histogram is updated at the tick-frequency of the incoming time series and the activity histogram is periodically updated, as described above.

Strictly speaking there are more periods than just one global DST period (summer time) and one non-DST period (winter time). In a market consisting of n components there are at most 2^n different periods. For a typical three-component market[4] this amounts to 8 periods. Since there is no DST in the East-Asian the number reduces to 4. During the last years there was at most a one-week lag between the DST switch in the US and in Europe, so that we can work, in a very good approximation, with two histograms only.

2.3 Market Components and Holiday Factors

If the seasonality of the market activity pattern were unperturbed no decompostion of the market into regional components would be needed.[5] The existence of public holidays, however, result in a perturbation of the seasonality. If a holiday applies for the whole market it can be taken into account by setting the market activity $a(t)$ to zero. Unfortunately there are only very few world-wide public holidays, while most of them are restricted to one or several regional markets, as can be seen from Tab. 1. Thus, it turned out to be necessary to decompose the market into conveniently chosen market regions. The reduction of market activity due to public holidays is then taken into account by setting to zero the activities $a_i(t)$ of the affected components, only.

The decomposition into market components and the holiday tracking mechanism are now described in turn.

[4]I.e., consisting of an American, an East-Asian, and a European component

[5]Even in that case, however, decomposition can prove useful to define a component-wise Θ Time, which makes it possible to decompose a given time series into several “partial” time series related to the regional components. Preliminary results of U. Müller show that interesting information about the capital flow behavior can be extracted from such a decomposition.

Date	Country and Name
01.01.1999	J: New Year's Day E: New Year's Day U: New Year's Day
15.01.1999	J: Adults Day
18.01.1999	U: MLK Day
11.02.1999	J: National Founding Day U: Lincolns Birthday
15.02.1999	U: Presidents Day
22.03.1999	J: Vernal Equinox Day
02.04.1999	E: Good Friday
05.04.1999	E: Easter Monday
29.04.1999	J: Greenery Day
03.05.1999	J: Constitution Day
04.05.1999	J: National Holiday
05.05.1999	J: Children's Day
31.05.1999	U: Memorial Day
05.07.1999	U: Independance Day Obs.
06.09.1999	U: Labor Day
15.09.1999	J: Respect F. T. Aged Day
23.09.1999	J: Autumnal Equinox Day
11.10.1999	J: Health-Sports Day U: Columbus Day
03.11.1999	J: Culture Day
11.11.1999	U: Veterans' Day
23.11.1999	J: Labor Thanksgiving Day
25.11.1999	U: Thanksgiving Day
23.12.1999	J: Emperor's Birthday
31.12.1999	J: New Years'Eve(Bank Hol.)

Date	Country and Name
17.01.2000	U: MLK Day
11.02.2000	J: National Founding Day U: Lincolns Birthday
21.02.2000	U: Presidents Day
21.03.2000	J: Vernal Equinox Day
21.04.2000	E: Good Friday
24.04.2000	E: Easter Monday
03.05.2000	J: Constitution Day
04.05.2000	J: National Holiday
05.05.2000	J: Children's Day
29.05.2000	U: Memorial Day
04.07.2000	U: Independance Day Obs.
04.09.2000	U: Labor Day
15.09.2000	J: Respect F. T. Aged Day
09.10.2000	U: Columbus Day
10.10.2000	J: Health-Sports Day
03.11.2000	J: Culture Day
23.11.2000	J: Labor Thanksgiving Day U: Thanksgiving Day
25.12.2000	E: Christmas U: Christmas
26.12.2000	E: Boxing Day

Table 1: Public holidays in 1999 and 2000 for Japan (J), the US (U), and major European countries[6] (E). Left column: Date. Right column: Market and name of holiday. Public holidays coinciding with a Saturday or Sunday are not listed. Notice that only the New Year's Day is observed in all three regions, the other holidays being restricted to at most two markets.

2.3.1 Regional Market Components

Decomposition of the total market activity into component activity is by no means easy. The main problem is that information about the origin of ticks cannot be used because it is not always available and tick activity is a very error-prone measure. Therefore, the decomposition into market components only relies on the fact that different components have different daily opening periods and some additional simplifying assumptions.

First, dynamic activity share factors $s_i(t)$ with $i = 0 \ldots n$ are introduced. For $i > 0$ they determine the instantaneous activity share of market component i, and the constant share s_0 accounts for the part of the activity not taken into account by the explicitly modelled components.The time-dependent factors $s_i(t)$ are products of the time-independent weight factors w_i indicating the relative weights of the respective market components and the opening functions $o_i(t)$, which is a kind of smoothed indicator function with value $o_i(t) \approx 1$ during the component's active period and $o_i(t) \approx 0$ during its

[6] Those relevant for FX trading.

Date	Country and Name	Time Period	Factor
15.09.1999	J: Respect F. T. Aged Day	whole day	0.3
hline 24.12.1999	E: Christmas' Eve	whole day	0.5
	U: Christmas' Eve	00:00:00–12:00:00	0.5
	U: Christmas' Eve	12:00:00–24:00:00	0.1
15.09.2000	J: Respect F. T. Aged Day	whole day	0.3
31.12.1999	E: New Years' Eve	whole day	0.5
	U: New Years' Eve	00:00:00–12:00:00	0.5
	U: New Years' Eve	12:00:00–24:00:00	0.1

Table 2: Fuzzy and partial holidays in 1999 and 2000 for Japan (J), the US (U), and major European countries (E). 1st column: Date. 2nd column: Indication of countries and name of day. 3rd column: Time period for partial holiday (in local time). 4th column: Value of holiday factor $h_i^{(\alpha)}$.

inactive period. The exact definition of the share factors is[7]

$$s_i(t) = \begin{cases} \dfrac{w_0}{w_0 + \sum\limits_{i=1}^{n} w_i o_i(t_i)} & \text{if } i = 0 \\[2ex] \dfrac{w_i o_i(t_i)}{w_0 + \sum\limits_{i=1}^{n} w_i o_i(t_i)} & \text{otherwise,} \end{cases} \tag{24}$$

where $t_i = t_i(t)$ is the local time of component i and w_0 is the background activity. The s_i observe the normalization condition

$$\sum_{i=0}^{n} s_i(t_i) = 1 \qquad \forall t_i. \tag{25}$$

The opening function $o_i(t)$ indicating the activity state is defined as

$$o_i(t) = \left(\frac{1}{1 + \exp\left\{-\gamma_i^{(o)}(t - t_o + \Delta)\right\}} \right) \left(\frac{1}{1 + \exp\left\{-\gamma_i^{(c)}(t_c - \Delta - t)\right\}} \right) \tag{26}$$

where the $\gamma_i^{(o/s)}$ are steepness parameters controlling the duration of the transition period, t_o (t_c) indicate the begin (end) of the activity period, and Δ_o (Δ_s) are shift parameters introduced to make t_o and t_c compatible with the static ϑ time components configuration.

The number of components can be chosen by the user. A three-component market with an American, an East-Asian, and a European component, which is the hard-wired setting for the static ϑ time, is still considered as default. Fig. 5 shows the component-wise activity or this setting. However, for localized, exchange-traded markets as the DAX, only one component is needed. On the other hand, a fourth component has been introduced for the USD/JPY spot rate to account for the influence of the Australian market before Tokyo opening. Other configurations may become necessary with the emergence of new markets.

2.3.2 Holiday Tracking Mechanism

Public holidays are taken into account by setting the activity of the affected components to zero.

[7] In the prototype, the share factors are computed from the polynomial model for the market activity by

$$s_i(t) = \frac{\alpha_i(t)}{\sum_{i=0}^{n} \alpha_i}$$

where the $\alpha_i(t)$ are the activities given by the polynomial model of the static θ time. However, since these factors have an effect only for periods of the day when the activity of market components overlaps, a single set of parameters for every type of instruments is enough in most cases.

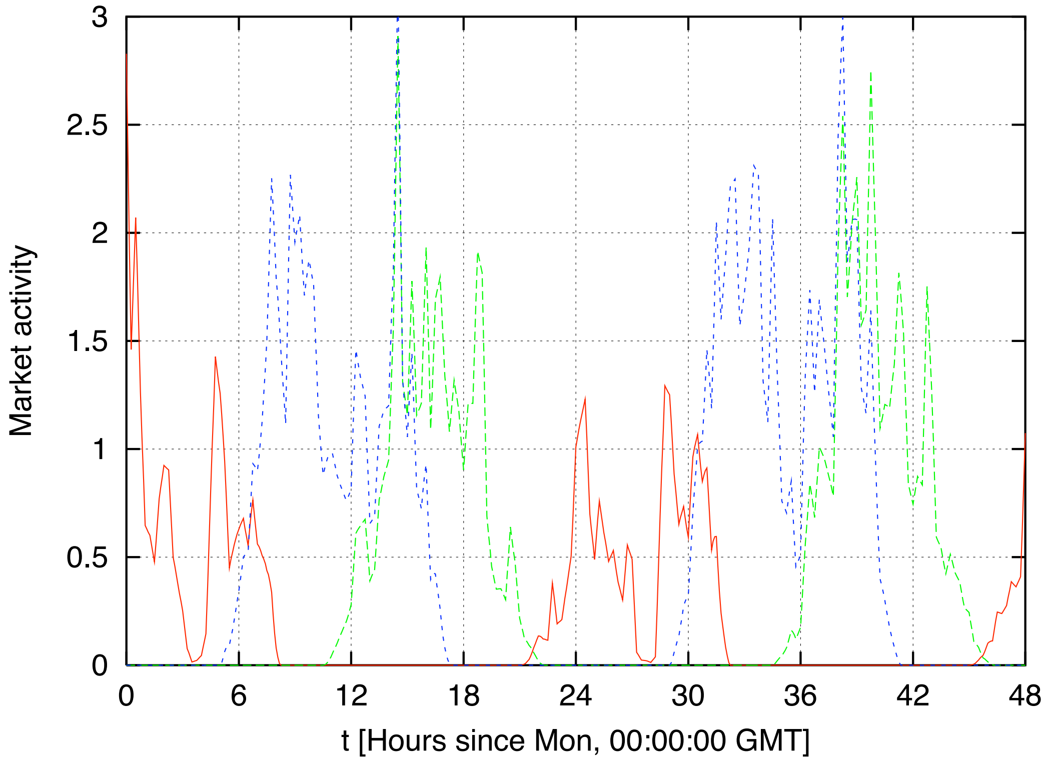

Figure 5: Activities of regional market components for a major FX spot rate. The market is decomposed into an East-Asian (red), an European (blue), and an American (green) component. For better visibility the activity is shown only for every other day.

In addition to public holidays, reduction of market activity is regularly observed during some other special days as, e.g., on Christmas' Eve or on Friday if preceeding Thursday is a public holiday. To account for these effects, the concepts of partial and fuzzy holidays are introduced. During a partial holiday the market is inactive only for a part of the normal activity period, while fuzzy holidays are days with reduced but still positive activity level.

The effect of these kind of exceptional days is taken into account by component specific holiday activity factors $h_i^{(\alpha)} \in [0,1]$ defined as

$$h_i(t) = \begin{cases} h_i^{(\alpha)} & \text{if } t \in [T_{i,b}^{(\alpha)}, T_{i,e}^{(\alpha)}] \\ 1 & \text{otherwise.} \end{cases} \tag{27}$$

Here, i relates to a specific component and α labels the public holidays. The holiday activity factor $h_i^{(\alpha)}$ is a positive number strictly smaller than 1, and $T_{i,b}^{(\alpha)}$ ($T_{i,b}^{(\alpha)}$) is the beginning (end) of holiday α of component i. For public holidays, $h_i^{(\alpha)} = 0$, $T_{i,b}^{(\alpha)} =$ 00:00:00 and $T_{i,e}^{(\alpha)} =$ 24:00:00. In the case of fuzzy holidays, $h_i^{(\alpha)} > 0$, and for partial holidays, $T_{i,b}^{(\alpha)} >$ 00:00:00 or $T_{i,e}^{(\alpha)} <$ 24:00:00. Examples of partial and fuzzy holidays are listed in Tab. 2.

3 Differences Between the Dynamic Θ Time and the OA Static ϑ Time

The Dynamic Θ Time is a substantial improvement over the static ϑ time developed by Olsen & Associates in 1990 and described in (Dacorogna *et al.*, 1993). The essential new features are the following.

- Better modelling of abrupt temporal changes in the weekly market activity pattern, and modelling day-of-the-week effects.
- Adaptability to slow changes in the market acitvity pattern over time.
- Extension of the spectrum of markets that can be treated. In particular, the Dynamic Θ Time makes it possible to deseasonalize the time series of localized, exchange traded instruments.

In addition, there are the following technical improvements.

- Fitting of the activity pattern is no longer needed.
- The treatments of holidays (partial and fuzzy holidays) and of DST periods has been improved.
- There is now the possibility to use quantities other than the volatility as proxy for market activity. Tests have been performed, e.g., with the tick rate (cf. Fig. 7) to gain information about the actual activity period. However, up to now no deseasonalization has been done using this quantity.

Technically these improvements have been achieved by the following measures.

- Separation of the histogram part, the regional component part and the holiday and DST treatment (cf. Fig. 6). This makes the modelling of the regional market components more flexible. In particular, the user can now choose any number of market components, and can easily configure the system for new components. This is possible because in contrast to the static ϑ time, where the sum of the component's polynomials represents the *absolute* activity pattern, in the case of the Dynamic Θ Time the regional component modelling has to account only for the components' time-depending relative weightings. These parameters vary much less from one instrument to the other and do not need to be configured for any instrument individually.
- Decomposition into regional market component is achieved by dynamic weight factors $w_i(t)$ instead of fitting with 7th and 5th order polynomials.
- The activity pattern is no longer modelled by 7th order polynomials, with require rather involved fitting and fail to model abrupt changes in the activity pattern, but by means of histograms whose resolution can be configured. The changes in the activity pattern are tracked using IWMAs.
- While the static ϑ time uses a static activity pattern, in the case of the Dynamic Θ Time the seasonal volatility pattern is computed from the incoming ticks by an IWMA. This assures adaptability of the pattern. The information newly acquired during the computing process is migrated to the activity pattern at regular time intervals.
- Separate histograms are used for the activity patterns during DST periods and no-DST periods, while in the static ϑ time the polynomials where fitted to a *mean* activity pattern averaged over both, DST and no-DST periods.
- The activity pattern update mechanism can easily be modified to take into account other kind of information as, e.g., tick rates or spreads.

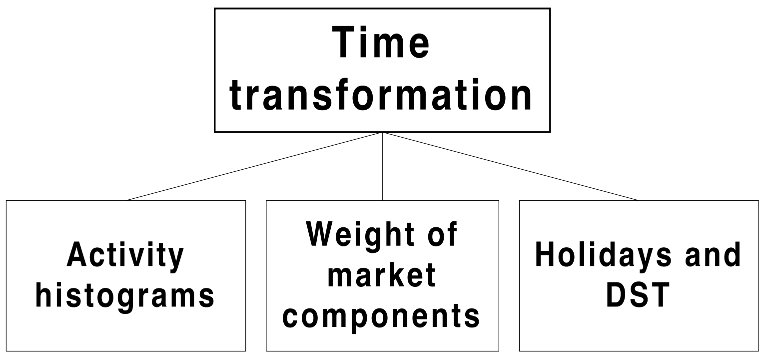

Figure 6: Components of the algorithm.

4 Two Examples

Here three configurations of the Dynamic Θ Time are presented, namely for typical FX spot rates, for USD/JPY spot rates, and for the the German DAX. While the main configuration parameters for these examples are displayed in Tab. 3, an account of the whole set of configuration parameters is relegated to App. C.

4.1 Dynamic Θ Time for FX Spot Rates

FX spot markets are over-the-counter markets, which in principle are active 24 hours a day seven days a week. This is due to the fact that the market is composed of different regional components whose activity periods overlap. Even though the market activity of a given regional component is zero during nights and weekends, there are no well-defined opening and closing hours.

As a whole, the volatility pattern is rather smooth (cf. Figs. 1 and 2). From Monday to Friday the activity is strongly reduced only between American closing and Japanese opening and during Japanese lunch time. Market activity falls to zero only during the weekend.

According to the main activity regions the market is decomposed into three regional components: the American, East-Asian, and European component. Careful studies showed that in the case of the USD/JPY spot rate there is some activity before the regular Tokyo opening, and this activity could be attributed to the Australian market activity. Accordingly, an Australian component has been added for this instrument.

The main configuration parameters are displayed in columns two (general FX spot rates) and three (USD/JPY) of Tab. 3. In addition to the regional market components, there is the configuration of the resolution of the histogram (bin width Δ), the range τ_s of the MA responsible for a first, local smoothing according to eq. (20) and the time step δ for the regularization of the incomming time series. The values of the parameters are adjusted to the relative smoothness of the volatility pattern.

4.2 Dynamic Θ Time for Equity Indices

Equity indices are instruments traded in a localized stock exchange. Since they are typically traded in a single place during well-defined opening hours, their volatility patterns show a sharp increase at market opening, a sharp drop to zero at market closing, and the well-known U-shaped pattern during the opening period.

Since the market is geographically localized, only a single market component is needed for the Dynamic Θ Time configuration (cf. 4th column in Tab. 3). The sharp drops at opening and closing require a good time resolution of the activity histogram. Therefore the bin width has been reduced to 15 minutes, and the other two time ranges have been shortened accordingly. Notice, however, the trade-off between resolution and noise: the higher the resolution the bigger the random noise.

The opening and closing time have changed during the period for which the results presented in the next section have been calculated. In 1996 the daily opening period was 3.5 hours (from 10:00 till 13:30 local time. This time has been extended (partly due to the introduction of computer trading)

Parameter	General FX	UDS/JPY	DAX
Market components (with relative weights)	American (1/3) East-Asian (1/3) European (1/3)	American (0.3125) East-Asian (0.3125) European (0.3125) Australian (0.0625)	German (1)
Bin width Δ of activity histograms	1 h	1 h	15 min
Range of short MA τ_s	30 min	30 min	7 min 30 sec
Sampling time unit δ	5 min	5 min	3 min 45 sec

Table 3: Comparison of the configuration for different instruments. The parameters Δ, τ_s, and δ are explained in Sec. 2.2.3.

to a period of about 8 hours in 1998. The Dynamic Θ Time was able to adapt to this change in the daily opening period smoothly, as is corroborated by the fact that the strong seasonal peaks in the volatility autocorrelation function have been removed nearly completely, as shown in Sec. 9 below.

5 Test of the Deseasonalization Quality

Two tests of the deseasonalization quality of the Dynamic Θ Time are presented: a measure of the residual seasonal fluctuations of the volatility after the time transformation and autocorrelation functions of the absolute returns.

5.1 Residual Seasonalities

The idea of this measure is the following: The deseasonalization algorithm should remove the seasonal part of the time series' volatility fluctuations. Thus, to measure the deseasonalization quality, a quantity is needed that is sensitive only to the seasonal flcutuations in the volatility and *not* to its random fluctuations.

A measure of the total volatility fluctuations is provided by the volatility of the volatility. To make this quantity time dependent it is defined in terms of moving averages as

$$\Delta v[\Delta\Theta](t) = \sqrt{\mathrm{MA}[v^2, T](t) - (\mathrm{MA}[v, T](t))^2}, \qquad (1)$$

Here $v \equiv v[\Delta\Theta](t)$ is the annualized volatility defined as 2-norm of the fluctuations with respect to a time horizon of $\Delta\Theta = 7\,\mathrm{min}\,30\,\mathrm{sec}$:

$$v[\Delta\Theta](t) = \sqrt{\mathrm{MA}\left[(r_{\mathrm{ann}}[\Delta\Theta](t))^2, \delta\right]} \qquad (2)$$

with the annualized return

$$r_{\mathrm{ann}}[\Delta\Theta](t) = (x(t) - x(t - \Delta\Theta)) \cdot \sqrt{\frac{1\,\mathrm{year}}{\Delta t(\Delta\Theta, t)}} \qquad (3)$$

and an MA with range $\tau = 30\,\mathrm{min}$.

Eq. (1) provides a measure for the total fluctuations of the volatility, which comprises the stochastic as well as the seasonal fluctuations. The latter can be separated from the former if intra-week statistics is used. Let v_i, $i = 1 \ldots N$ be the values of a weekly volatility histogram computed by an IWMA with

Name	phycial time		Dyn. Θ Time		static ϑ time		improvement[%]	
	$\langle\Delta v_{\mathrm{seas}}\rangle_t$ $\times10^{-2}$	$\langle\delta(\Delta v_{\mathrm{seas}})\rangle_t$ $\times10^{-3}$	$\langle\Delta v_{\mathrm{seas}}\rangle_t$ $\times10^{-2}$	$\langle\delta(\Delta v_{\mathrm{seas}})\rangle_t$ $\times10^{-3}$	$\langle\Delta v_{\mathrm{seas}}\rangle_t$ $\times10^{-2}$	$\langle\delta(\Delta v_{\mathrm{seas}})\rangle_t$ $\times10^{-3}$	i_1	i_2
EUR/JPY	3.0	4.3	1.1	1.6	1.6	5.0	47.2	216.2
EUR/USD	2.7	4.7	0.7	1.7	1.1	2.2	65.3	29.1
GBP/DEM	2.7	5.8	0.8	2.7	1.6	3.8	93.1	40.2
GBP/JPY	4.7	9.3	1.9	4.2	6.1	11.8	226.6	176.9
GBP/USD	3.6	5.4	0.9	1.7	1.2	2.0	26.9	19.5
USD/ATS	4.1	4.9	1.3	2.6	3.8	10.9	183.4	316.8
USD/CAD	1.9	4.4	0.5	0.6	0.7	1.6	56.0	164.7
USD/DEM	3.6	7.2	0.8	1.8	1.3	2.5	48.3	38.0

Table 4: Deseasonalization quality for 8 major FX spot rates. $\langle\Delta v_{\mathrm{seas}}\rangle_t$ is the mean of the values of Δv_{seas} defined in eq. (4) reported weekly during the 5-year period from 01.01.1994 till 31.12.1998, which are displayed in Fig. 8. The quantities i_1 and i_2 are defined in eqs. (5) and (6), repsectively. A 15-month build-up starting around 01.10.1992 has been used to initialize the Dynamic Θ Time as well as the various moving averages required for data analysis.

a range of 4 month, as described in section 2.2.3. Then, the seasonal part of the fluctuation, i.e., the *histogram volatility* can be defined as

$$\Delta v_{\text{seas}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} v_i^2 - \left(\frac{1}{N}\sum_{i=1}^{N} v_i\right)^2} \tag{4}$$

where N is the number of bins in the intra-week histogram.

This measure does precisely what is needed: It acts as a kind of "low pass filter" with a cut-off frequency of the inverse of the width of the histogram bins. Thus, it only reports the part of the volatility fluctuation contained in the fluctuations of the values of the bins, which essentially are the pattern's seasonal fluctuations.

In order to take into accounts the weekends, an appropriate weighting was introduced. Setting the weight to zero for the period from Friday, 21:00 GMT till Sunday, 21:00 GMT turned out to work well.

Fig. 8 displays the time dependence of the ratios $\Delta v_{\text{seas}}^{(\text{static }\vartheta)}/\Delta v_{\text{seas}}^{(\text{physical})}$ (top) and $\Delta v_{\text{seas}}^{(\text{Dyn. }\Theta)}/\Delta v_{\text{seas}}^{(\text{physical})}$ (bottom) for a period of 5 years. The reduction of the remaining seasonality and the gain in stability is striking and do not need any further comment. Averaging over the whole 5 year period yields the quantitative values displayed in Table 4. The last two columns show the reduction in % defined as:

$$i_1 = 100\left(\frac{\langle \Delta v_{\text{seas}}^{(\text{static }\vartheta)}\rangle_t}{\langle \Delta v_{\text{seas}}^{(\text{Dyn. }\Theta)}\rangle_t} - 1\right) \tag{5}$$

for the mean and as

$$i_2 = 100\left(\frac{\langle \delta(\Delta v_{\text{seas}}^{(\text{static }\vartheta)})\rangle_t}{\langle \delta(\Delta v_{\text{seas}}^{(\text{Dyn. }\Theta)})\rangle_t} - 1\right) \tag{6}$$

for the variations of the remaining fluctuations. The $\langle \ldots \rangle_t$ indicates the temporal average over the 5-year testing period. The remaining volatility fluctuations are reduced by up to more than 200% (GBP/JPY) and their fluctuations are reduced by up to more than 300% (USD/ATS).

5.2 Volatility Autocorrelation Functions

Volatility autocorrelation functions provide another very sensitive test of the deseasonalization quality. The autocorrelation function $c[x](\tau)$ of a time series $x(t)$ is defined as

$$c[x](\tau) = \frac{\langle (x(t) - \langle x \rangle)(x(t+\tau) - \langle x \rangle)\rangle_t}{\sqrt{\langle x^2(t)\rangle_t \langle x^2(t+\tau)\rangle_t}}. \tag{7}$$

In the present case, $x(t) = v[\Delta t](t)$ with $\Delta t = 1\,\text{hour}$ and the 1-norm of the return is chosen to measure the volatility. The result is displayed in Fig. 9 for USD/JPY spot rates (left) and the German DAX (right). In both cases there are very pronounced seasonalities with one day and one week period in physical time (green lines). In Dynamic Θ Time (red lines), these seasonalities are completely removed for the FX rate (left, red line). This is in line with the results of the previous section. The corresponding figure computed for the same quantity in static ϑ time shows a less perfect deseasonalization. For the DAX (right) deseasonalization is also very good, even though the daily and weekly periodicity is still distinguishable. Since the DAX is a localized exchange-traded instrument, no deseasonalization is available with the static ϑ time.

A Exponential Moving Averages

A.1 EMA, Iterated EMA, and Moving Average

This paragraph is a summary of those results of sections 4.1–4.3 of (Zumbach and Müller, 2001) on which the θ-time algorithm relies. For more detailed information it is refered to that document.

A.1.1 Exponential Moving Average EMA$[\tau]$

The basic exponential moving average (EMA) is a linear operator defined by the convolution

$$\text{EMA}[\tau; z](t) \quad = \quad \int_{-\infty}^{t} dt' \, \text{ema}[\tau](t - t') \, z(t') \tag{A.1}$$

with the exponentially decaying kernel

$$\text{ema}[\tau](t) = \frac{e^{-t/\tau}}{\tau} \ . \tag{A.2}$$

The numerical evaluation can efficiently be done by the simple iterative formula (Müller, 1991):

$$\text{EMA}[\tau; z](t_n) \quad = \quad \mu \ \text{EMA}[\tau; z](t_{n-1}) + (\nu - \mu) \, z_{n-1} + (1 - \nu) \, z_n \tag{A.3}$$

with

$$\begin{aligned} \alpha &= \frac{t_n - t_{n-1}}{\tau} \\ \mu &= e^{-\alpha} \end{aligned}$$

and ν depends on the chosen interpolation scheme,

$$\nu = \begin{cases} 1 & \text{previous point} \\ (1 - \mu)/\alpha & \text{linear interpolation} \\ \mu & \text{next point.} \end{cases} \tag{A.4}$$

A.1.2 The Iterated EMA$[\tau, n]$

Iterating the basic EMA operator provides a family of iterated exponential moving averages EMA$[\tau, n]$, which can recursively be defined as

$$\text{EMA}[\tau, n; z] = \text{EMA}[\tau; \text{EMA}[\tau, n - 1; z]] \tag{A.5}$$

with EMA$[\tau, 1; z] \equiv$ EMA$[\tau; z]$. Iteratived EMAs can efficiently be evaluated by means of the iteration (A.3) for all its basic EMAs. The kernel of EMA$[\tau, n]$ is

$$\text{ema}[\tau, n](t) = \frac{1}{(n-1)!} \left(\frac{t}{\tau} \right)^{n-1} \frac{e^{-t/\tau}}{\tau} \ . \tag{A.6}$$

The range, width and aspect ratio of the iterated EMA are

$$\begin{aligned} r &= n\tau \ , \\ \langle t^2 \rangle &= n(n+1)\tau^2 \ , \\ w^2 &= n\tau^2 \ , \\ AR &= \sqrt{(n+1)/n} \ . \end{aligned} \tag{A.7}$$

A.1.3 Moving Average MA$[\tau, n]$

A moving average (MA) is an integral operator with a retangular kernel:

$$\text{MA}[\tau, z] = \int_{-\infty}^{t} \text{d}t' \text{ma}[\tau](t - t') \, z(t') \tag{A.8}$$

with

$$\text{ma}[\tau](t) = \begin{cases} \frac{1}{2\tau} & \text{if } 0 \leq t < 2\tau \\ 0 & \text{elsewhere.} \end{cases} \tag{A.9}$$

The MA can then be reexpressed as a superposition of EMAs:

$$\text{MA}[\tau, n] = \frac{1}{n} \sum_{k=1}^{n} \text{EMA}[\tau', k] \qquad \text{with } \tau' = \frac{2\tau}{\text{n} + 1} \ . \tag{A.10}$$

The choice of the parameter τ' assures that the range of MA$[\tau, n]$ is $r = \tau$, independently of n. This provides a family of more and more rectangular-shaped kernels, with the relative weight of the distant past controlled by n. The analytical form of the kernels is given by

$$\text{ma}[\tau, n](t) \quad = \quad \frac{1}{n} \sum_{k=1}^{n} \text{ema}[\tau, n](t) \tag{A.11}$$

with $\text{ema}[\tau, n](t)$ defined in eq. (A.6). For large n, this operator tends to a rectangular moving average. For n values of $n \sim 5$ and higher, the kernel is rectangular-like more than EMA-like but still tends to zero smoothly rather than abruptly. This has the advantage of avoiding spurios noise, unlike the rectangular kernel of the exact moving average. The aspect ratio of the MA operator is

$$AR = \sqrt{\frac{4(n+2)}{3(n+1)}} \ . \tag{A.12}$$

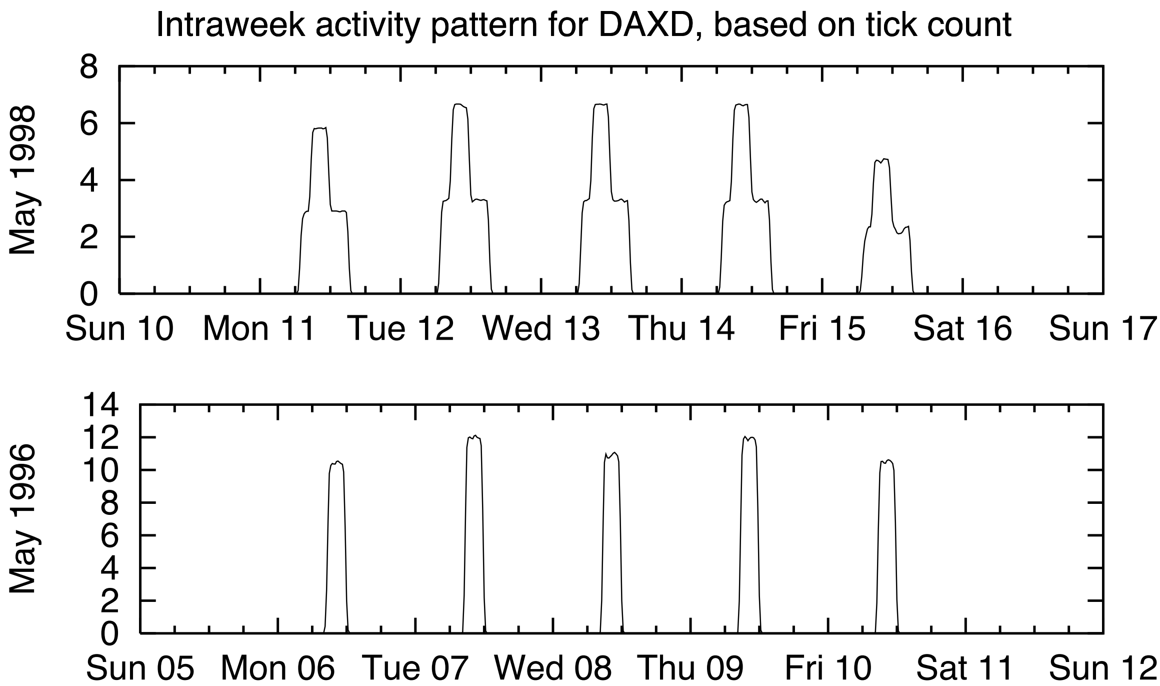

Figure 7: Activity period of the German DAX in May 1996 (bottom) and in May 1998 (top) based on tick counts. Notice the jump in the tick rate at market opening and closing, and the extended activity periods in 1998. The averaging period for the autocorrelation function displayed in Fig. 9 extends over this change in the activity period, proving the efficiency of the adaption mechanism of the Dynamic Θ Time.

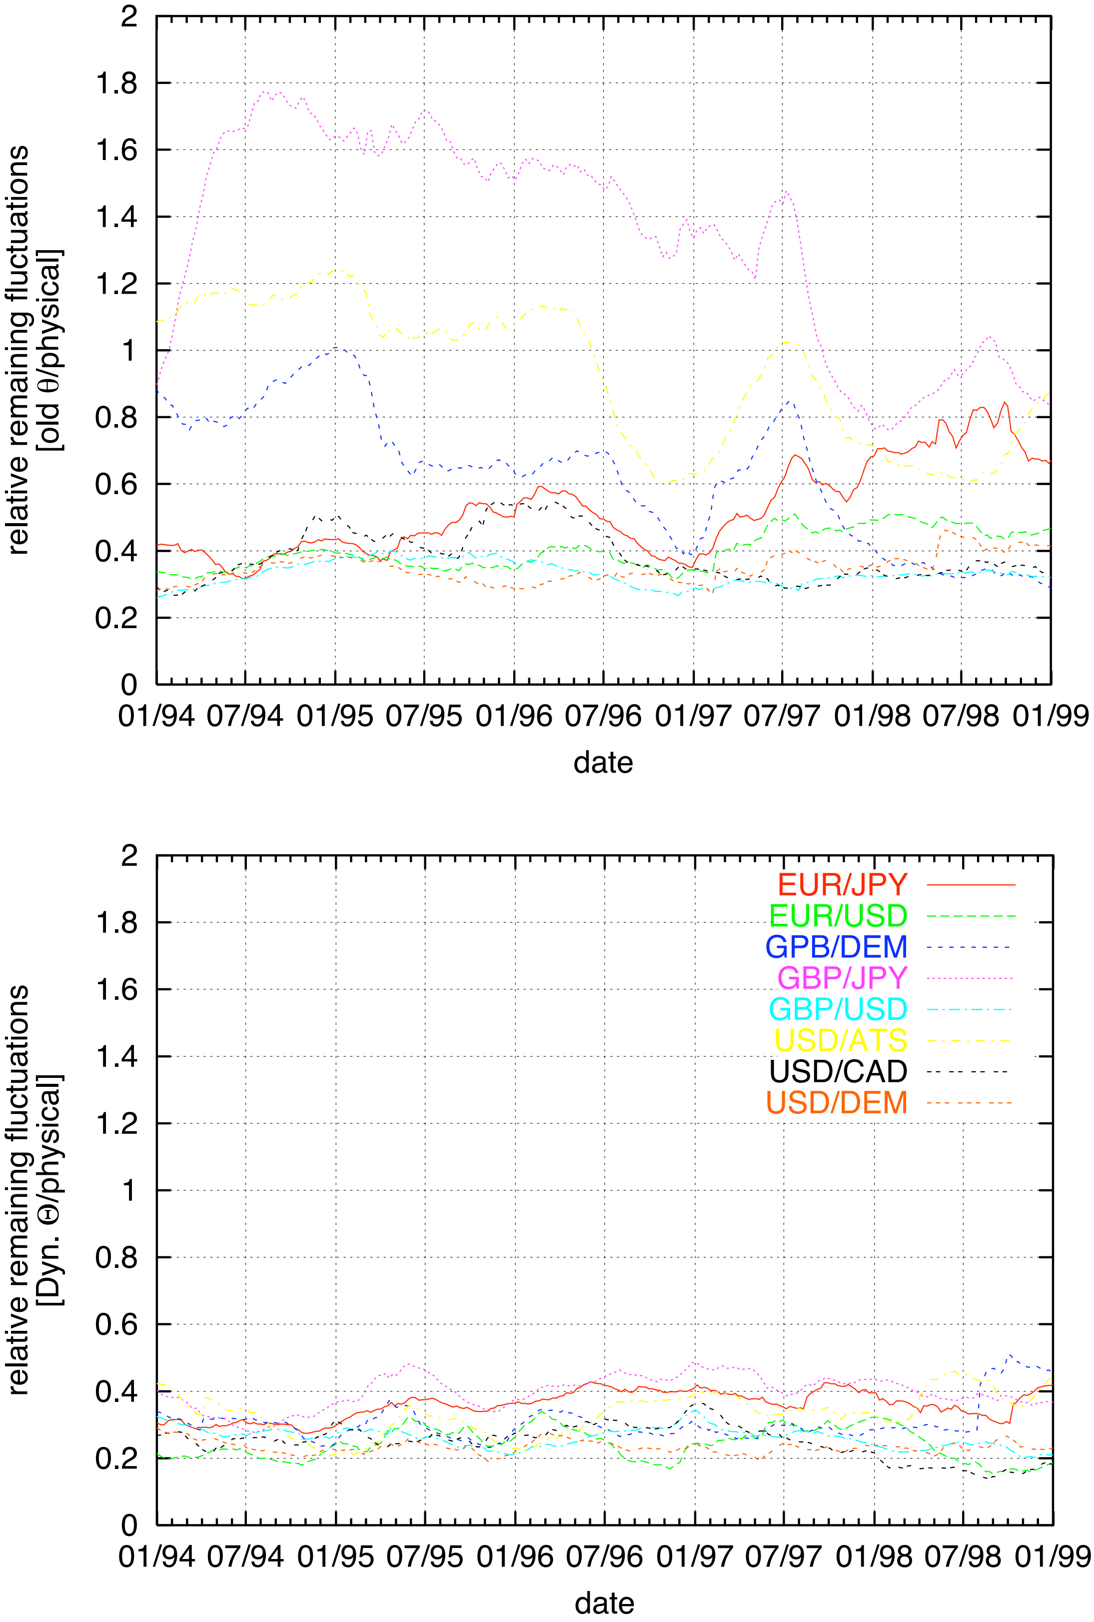

Figure 8: Ratio of seasonal fluctuations in physcial time and remaining residual seasonal fluctuations, $\Delta v_{\mathrm{seas}}^{(\mathrm{static}\ \vartheta)}/\Delta v_{\mathrm{seas}}^{(\mathrm{physical})}$ (top) and $\Delta v_{\mathrm{seas}}^{(\mathrm{Dyn.}\ \Theta)}/\Delta v_{\mathrm{seas}}^{(\mathrm{physical})}$ (bottom) , after transformation to static ϑ time and Dynamic Θ Time, respectively. Eq. (4) is used to compute the seasonal fluctuation measure v_{seas}.

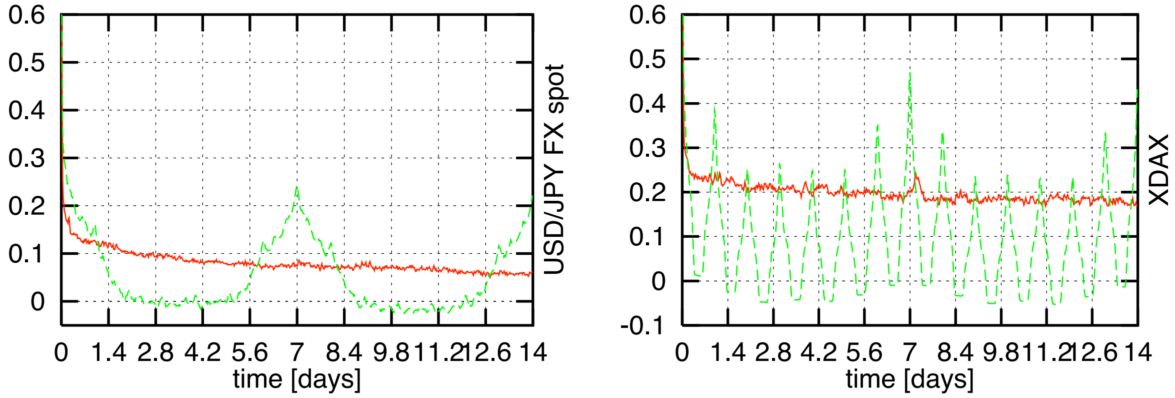

Figure 9: ACF for (a) USD/JPY (1/1990–12/1999) and (b) German DAX (1/1996–12/1998). Full red lines: Dynamic Θ Time. Dashed green lines: Physical time scale. Notice that in Dynamic Θ Time the daily seasonalities produce peaks with a periodicity of approximately 1.4 days (1 Θ-day). No residual seasonalities can be detected for the FX rate while very week residual daily seasonalities and somewhat larger weekly seasonalities remain in the case of the DAX.

B Summary of the Algorithm

A summary of the algorithm is displayed in the following table.

	Eq. no.	Equation
$\Delta\Theta \leftarrow$ Market activity	(1)	$\Theta(t_2) - \Theta(t_1) = \Delta\Theta(t_1, t_2) = \int_{t_1}^{t_2} a(t)dt$
Decomposition into components	(3)	$a(t) = a_0 + \sum_{i=1}^{n} a_i(t),$
Single component activity	(4)	$a_i(t) = \left(\frac{s_0}{n} + h_i(t)\, s_i(t)\right) \text{IWMA}[\tau; a_\alpha](t),$
Market activity $\leftarrow$ volatility	(9)	$\text{IWMA}[\tau; a](t) = c \cdot \text{IWMA}[\tau; v[\Delta T]]^2(t)$
Volatility smoothing	(20)	$v_{\text{sm}}(t_i) = \frac{1}{\delta} \int_{t_i - \delta}^{t_i + \delta} \omega(t' - t_i)\ \text{MA}[\tau_s, v_{\text{reg}}[\Delta T]](t')\, \text{d}t'$
Definition of volatility IWMA	(22)	$\text{IWMA}[\tau, m; v](t_i) \equiv \text{IWMA}[\tau, m; v_{\text{sm}}](t_i) = \sum_{n=0}^{\infty} v_{\text{sm}}(t_i - nT_w)\ \text{iwma}[\tau, m]\,(n)$
Update of volatility IWMA	(23)	$\text{IWEMA}[\tau', k; v_{\text{sm}}]\,(t_i) = \begin{cases} \mu\ \text{IWEMA}[\tau', k; v_{\text{sm}}]\,(t_i - T_w) + (1-\mu)\, v_{\text{sm}}(t_i) & \text{if } k = 1 \\ \mu\ \text{IWEMA}[\tau', k; v_{\text{sm}}]\,((t_i - T_w) + (1-\mu)\,\text{IWEMA}[\tau', k-1; v_{\text{sm}}]\,(t_i) & \text{otherwise,} \end{cases}$
Share factors	(24)	$s_i(t) = \frac{w_i o_i(t_i)}{w_0 + \sum_{i=1}^{n} w_i o_i(t_i)}$
Opening function	(26)	$o_i(t) = \frac{1}{\left(1+\exp\left\{-\gamma_i^{(o)}(t-t_o+\Delta)\right\}\right)\left(1+\exp\left\{-\gamma_i^{(c)}(t_c-\Delta-t)\right\}\right)}$
Holiday factors	(27)	$h_i(t) = \begin{cases} h_i^{(\alpha)} & \text{if } t \in [T_{i,b}^{(\alpha)}, T_{i,e}^{(\alpha)}] \\ 1 & \text{otherwise.} \end{cases}$

Table 5: Summary of the Dynamic Θ Time algorithm

C Configuration

C.1 Global Parameters

Top Level Configuration Keys

```
( scalingLawExponentD      0.5   )
( scalingLawExponentC    -11.990 )
```

They provide the scaling exponent and the offset for the scaling law of the mean absolute value of the returns, as defined in (Müller *et al.*, 1990).

```
( marketWeights
    (...)
)
```

Contains the following subkeys for the configuration of the decomposition into regional market components.

- Model for computing the dynamic share factors $s_i(t_i)$:

  ```
  ( weightModel       stepFunctionWithLocalTime)
  ```

- Constant weight w_0, as used in eq. (24):

  ```
  ( constantComponent                      0.01)
  ```

- Configuration of the regional market components

  ```
  ( components                [(...) ...(...)] )
  ```

  The value is a list of config pairs whose structure is described in section C.2.

The top level key

```
( marketActivity
    (...)
)
```

contains the following subkeys:

- Time interval between two consecutive updates of the activity histogram from the corresponding volatility histogram:

  ```
  ("activityHistogramUpdateInterval"  "7d 00:00:00" )
  ```

- Controls whether Dynamic Θ Time starts with pre-computed histograms:

  ```
  ( noHistograms             true )
  ```

- File names of pre-computed histograms if any:

  ```
  ( activityHistograms      [...] )
  ```

- Value of the constant activity a_0:

  ```
  ("constantActivity"       0.001 )
  ```

- Controls the use of the hole-detection mechanism:

  ```
  ("holeDetection"          false )
  ```

The top-level key

```
(ObcIntraWeekMA
    (...)
)
```

contains the following subkeys:

- The number of intervals in the intra-week histograms:

```
( nbIntervals              168 )
```

- The center of gravity of the IWMA, i.e., half of the width of the MA window:

```
("MA center of gravity"   "30d 00:00:00" )
```

- The order m of the approximation of the rectangular MA window, as defined in eq. (19):

```
("orderMax"               8 )
```

- Configuration of the regularizer

```
( ObcRTkSDifference
    (...)
)
```

  with the following regularizer configuration keys:
  - Time step δ of the regularized volatility series refered to in eq. (20):
    ```
    ("sampling time interval"    "00:05:00")
    ```
  - Reference time unit for standardization of volatility:
    ```
    ("time unit"             "365d 06:00:00") # 1 year
    ```
  - Time horizon of volatility, ΔT = (number of steps) $\times\ \delta$:
    ```
    ("number of steps"                   2 )
    ```
  - Control of scaling during data gaps:
    ```
    ("diffusionLikeGapScaling"     "false" )
    ```

- Configuration of τ_s used in eq. (20)) for smoothing of the volatility series:

```
("shortMARange"              "00:30:00")
```

The values refer to the default configuration.

C.2 Market Component Configuration

The key `constantComponent` takes as values a list of component-specific config-pairs. Each of them contains the following entries:

- The name of the regional market component:

```
( marketName  EastAsian)
```

- Market opening or begin of the daily activity period[8] in local time:

```
( opening   "06:00:00" )
```

- Market closing or end of daily activity period in local time:

```
( closing   "17:15:00" )
```

  Weight of market component, as used in eq. (24):

```
( weight            1.)
```

- Opening slope of activity indicator function:

```
( openingSlope       4.)     #   in hour^-1
```

- Closing slope of activity indicator function:

```
( closingSlope       4.)     #   in hour^-1
```

- Shifts of opening and closing time with respect to the center of the opening and closing flank (parameters Δ_o and Δ_c in eq. (26)):

```
( shift     "01:30:00" )
```

The values listed above refer to an East-Asian market component.

[8]Since there is no well-defined opening time in OTC markets, we prefer to use the term *daily activity period* in that case.

References

Breymann W., 1998, *Project Design – Revision of Business time*, Internal document WAB.1998-05-04, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

Breymann W., Zumbach G., Dacorogna M. M., and Müller U. A., 2000, *Dynamical Deseasonalization in OTC and Localized Exchange-Traded Markets*, Internal document WAB.2000-01-31, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

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

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., Dacorogna M. M., Olsen R. B., Pictet O. V., Schwarz M., and Morgenegg C., 1990, *Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law, and Intraday Analysis*, Journal of Banking and Finance, **14**, 1189–1208.

Zumbach G. O. and Müller U. A., 2001, *Operators on Inhomogeneous Time Series*, International Journal of Theoretical and Applied Finance, **4**(1), 147–178.