Introducing a Scale of Market Shocks

GILLES O. ZUMBACH, MICHEL M. DACOROGNA,
JØRGEN L. OLSEN AND RICHARD B. OLSEN

GOZ.1998-10-01 (scaleOfMarketShocks)

March 23, 1999

A preprint by the O&A Research Group

Olsen & Associates
Research Institute for Applied Economics
Seefeldstrasse 233
8008 Zürich, Switzerland

Abstract

We introduce two "event" scales for financial markets, called "scale of market shocks" (SMS), which measure the importance of the market movements. These indices are based on the price volatility and are computed by integrating mapped asset volatilities over time horizons that range from 1 hour to 42 days. The first SMS is an absolute scale, or universal scale, allowing values of different assets to be compared directly. The second SMS is an adaptive scale, calibrated to the typical behavior of each asset, allowing the relative importance of market movements to be assessed. In principle, the SMS can be constructed for any market: the indices are computed from the price time series. In the foreign exchange (FX) market, each index is associated with a currency pair and we derive from it an index per currency and an index for the whole market.

In order to define the most appropriate SMS, the probability distribution for the volatilities is studied first. Then, the probability distribution of the two scales is computed. For USD/DEM and USD/JPY, the relations between peaks for the SMS and major "world events" is put forward. In addition, we also measure the correlation between the Scale of Market Shocks index and the size of the next price movements, which shows a high correlation for short time intervals.

Contents

1 Introduction

The financial markets often experience large price movements. Examples of extreme movements taken from the foreign exchange (FX) market are provided by the European Monetary System crisis of September 1992 or in October 1998 when the US Dollar lost 9% in one day against the Japanese Yen (October 7, 1998) and more than 18% in a week. One measurable consequence of these large and sudden price changes is the fat-tailed return distributions (Mandelbrot, 1963; Taylor, 1986; Hsieh, 1988; Müller et al., 1990), namely that the probability of a large price change decays as slow as a power law. When a market experiences such turbulence, it is common for the media to talk about a 'crash' or a 'crisis', but there is no serious attempt to quantify these events. In order to go beyond a heuristic analogy, it is essential to develop ways to measure the state of the financial market accurately and to provide quantitative assessment of market conditions. Such an objective assessment of the market state might help alleviate phenomena that often accompany 'crises' such as the widespread uncertainty over the reliability and stability of the financial system due to the strong interconnection between financial markets. An obvious illustration of a confidence crisis is provided by the recent events in Asia and Russia which culminated in August/September 1998.

Until now, there are no accepted ways of measuring *how large* the shocks are to which the market is submitted. However, in other fields, there are well accepted scales to measure the strength of an event or a shock. A familiar example is the Richter scale in geophysics (Richter, 1958). This scale is widely used and accepted to measure the intensity of an earthquake. In sailing, a similar scale exists, called the Beaufort scale, which measures the speed of winds and the state of the sea. The advantage of a well accepted scale is that shocks can be compared to each other and risk measures can be derived from them. Moreover, if the scale is well designed, it can serve as a warning signal that the market is entering a turbulent state.

The purpose of this paper is to introduce a similar 'event' scale in financial markets which would measure the importance of market movements. Such a measure provides a new analysis tool to the market, that is a new indicator. Due to the widely diverse characteristics of different financial assets, two scales are needed (they are presented in detail in section 5):

- a universal scale which allows scale values for different assets to be compared directly and
- an adaptive scale which is calibrated to the typical behavior of each individual asset.

These scales of market shocks, in short SMS, can be computed for any asset. In this paper, the emphasis of the empirical studies is on the foreign exchange market but the methodology can be applied with little modifications to other markets too. In order for this scale to provide a useful measure, it must be related to external events, or news, and calibrated on a wide range of these, ranging from 'average' market behavior to the most extreme cases.

We present here the specification and a first empirical study of the scale of market shocks. After having presented the basic idea in section 2, we introduce the formalism in sections 3. In section 4, we define the volatility and study the properties of its probability distribution. The definition of the SMS for a particular FX-rate is given and discussed in section 5 while in section 8 the definition is extended to obtain a scale for the whole market. In section 7, we compare the index S with the news headlines in order to make a connection with actual events. Then, in the context of risk forecasting, we measure the correlation of the scale of market shocks with the size of the next price movements in section 9. The conclusions are presented in section 10.

2 The basic idea

As a first guideline, we can proceed by analogy with Richter's approach in the construction of his famous scale (Richter, 1958). He defined a measure of the logarithm of the total energy liberated during an earthquake. As earthquakes are distributed with a power law probability distribution, the Richter scale also measures the (inverse) probability of an event. By analogy, we want to construct a logarithmic scale, namely a scale on which one more point corresponds to an event of the double intensity (more precisely, to a multiplication by a constant factor), or to an event which is twice as unlikely (more precisely, which is more unlikely by a constant factor). As we shall see below, these two definitions must be adapted for financial markets.

For the analogy with the Richter scale, we need to define the equivalent for financial markets of the concept of energy and total energy.

The energy: In mechanics, for a unit mass, the energy is given by $E = \vec{v}^2$ where $\vec{v} = d\vec{x}/dt$ is a speed, or a change of position per time interval. In finance, a possible analogy for the speed must be related to price changes. Price changes are created by an inbalance between buyers and sellers and correspond to the net flux of money. A price change is similar to a speed, i.e. to $Dx(t) = (x(t) - x(t-\tau)) \sim \vec{v}$, where x is the logarithmic middle price. However, with a stochastic process, we have one more parameter, namely the time range τ of the derivative $D[\tau]x$. Then, the (mechanical) energy E, at a given time range τ, is given by $E[\tau] \sim (D[\tau]x)^2$, i.e. the energy is related to a volatility measure. Therefore, we build the scale of market shocks from the instantaneous volatility, measured as an average over some time ranges of $\langle (D[\tau]x)^2 \rangle^{1/2}$.

If more information is available, for example exchanged volumes, it could also be incorporated in the definition of the energy. Volume information is unfortunately not known on the foreign exchange market since it is an over-the-counter market and there is no centralized place where transactions are recorded. The situation might change with the advent of automatic dealing systems. If this development becomes really significant and incorporates most of the market volume, another primary indicator could be used to build the SMS index, incorporating more information about the market.

The total energy: In an earthquake, the events are clearly separated: a beginning and an end can be identified, and the total energy released by an event can be integrated. The Richter scale corresponds then to the logarithm of the total energy. A financial process has a fundamentally different character, since it is dominated by the random component; a market is always fluctuating and moving, and this at all frequencies. Therefore, we cannot identify "events", with a clear beginning and end. The analogy with the Richter scale is limited here.

For a financial process, we want a continuous indicator, which is intuitively related to a total flux, or to the inbalance between buyers and sellers. In this sense, the mechanical work is a good physical analogy, since it is the rate of change of energy dE/dt.

The above discussion suggests to define the (instantaneous) scale of market shocks as the logarithm of the volatility at the time ranges τ, integrated over the different time ranges. Therefore, we take the following form for the scale of market shocks indicator S:

$$S[\mu, f; x] = \int d\ln\tau \;\; \mu(\ln\tau) \;\; f(v[\tau; x]) \tag{2.1}$$

where the function $f(\cdot)$ needs to be appropriately chosen. The measure $\mu(\ln \tau)$ fixes the weight given to the contribution at different time horizons τ. A part of this measure has been included in the term $d \ln \tau$, which reflects that the time range integral will be evaluated on a logarithmic scale. Formally, the limits of integration are from 0 to ∞, with some smooth cut-off included in the measure μ.

We want to construct an indicator sensitive to all market components, from intra-day dealers to long-term pension funds or central banks, that is why we need to integrate over τ. Clearly, these different market participants look at the same price curve with different time horizons. A real shock happens when all market participants become involved. The τ integration is essentially summing over the market components and the indicator S becomes large when the volatility is large at all time scales. In physics, this is similar to a second order phase transition. These transitions are dominated by fluctuations at all length scales, leading to diverging quantities like the specific heat or the magnetic susceptibility. Beside, our unpublished studies of volatilities defined at different time horizons show their relatively small correlations. Moreover, we know that there are asymmetries in the information flow between volatilities measured at different frequencies (Müller et al., 1997). These facts point to the relative independence and different information content of volatilities defined at different time horizons, and that there is no unique underlying volatility.

In order to turn this first form for S into a robust definition, we need to formalize the terms in eq. 2.1 precisely, namely

- the derivative and volatility operators,
- the form of the mapping function f
- and the measure μ.

As we are working with high frequency, real time data, which is not homogeneously spaced in time, a bit of sophistication is required here.

3 Notation and computation of the volatility

In this section, we present the notation, the basic definitions, and the main idea used for computing the volatilities. The exact definitions of the operators we are using here are described in more details in (Zumbach, 1998a).

Let us first fix the notation. Time series are denoted by a simple letter, like x. The value at time t of a time series x is denoted by $x(t)$. If a time series depends on a parameter p, it is denoted within brackets $x[p]$. If an indicator S is computed from another time series x, it is denoted by $S[p;x]$. For example, the scale of market shocks indicator $S[\mu;x]$ for a time series x depends on the measure μ. For a linear operator, we also use the notation $D[p;x] = D[p]x$, to make the linearity properties explicit.

Tick-by-tick (high frequency) data contains a time stamp t and the bid and ask prices p_{bid}, p_{ask} (Müller et al., 1990). We will consider the logarithmic middle price x as our primary time series

$$x = \frac{1}{2} \{\ln(p_{\text{bid}}) + \ln(p_{\text{ask}})\} \tag{3.2}$$

The annualized return in a time interval τ is then given by

$$r[\tau;x](t) = \frac{x(t) - x(t-\tau)}{\sqrt{\tau/1y}}. \tag{3.3}$$

The denominator is here to remove the Gaussian random walk scaling $\mathrm{E}\left[(x(t) - x(t-\tau))^2\right] \sim \tau$. With such a definition, $\mathrm{E}\left[r[\tau]^2\right] = \sigma^2$, with σ independent of τ to a very good approximation as we will see. The time interval τ needs to be expressed in some time unit. We choose years as units, denoted by the symbol $1y$ so that the returns are annualized.

The volatility v at the time scale τ of a time series x can be measured by

$$v[\tau;x](t) = \sqrt{\frac{1}{16}\sum_{i=1}^{16}\left(r[\tau/16;x](t_i)\right)^2} \tag{3.4}$$

with $t_i = t - (i-1)\tau/16$. This corresponds to an annualized volatility, because the return is already scaled by $1/\sqrt{\tau}$. As discussed in details in (Zumbach, 1998a), these last two formulae suffer from a number of drawbacks, particularly when working with high frequency data. Therefore, we instead measure the volatility with an equivalent formula based on exponential moving average (EMA) technology:

$$v[\tau;x] = \mathrm{MNorm}[\tau/2, p=2; D[\tau/16;x]] \tag{3.5}$$

The D operator computes a return similarly to 3.3, but using a smooth kernel built with an appropriate combination of EMA operators. The $\mathrm{MNorm}[\tau,p]$ operator computes a moving p norm in a window of length $\simeq 2\tau$. Usually, $p = 2$ is taken, but we will also use $p = 1/2$ below in order to have a more robust measure for the average volatility. For more details, the reader should refer to (Zumbach, 1998a).

A further problem is the treatment of daily and weekly seasonalities. This is a major issue as we are working in the intra day time range and because volatilities exhibit strong seasonalities (Müller et al., 1990; Baillie and Bollerslev, 1990). Without properly accounting for these effects, we would obtain a peak every European afternoon, corresponding simply to the trivial overlap of the European and American markets. For this reason, the above computations are done in the ϑ-time scale, as introduced by (Dacorogna et al., 1993). For more details, the reader should consult this reference.

4 The probability distribution for the volatility

It is generally assumed that the log-normal distribution is a good approximation for the probability distribution of volatility

$$p(v) = \frac{1}{\sqrt{2\pi}\sigma v}\exp\left[-\frac{1}{2\sigma^2}\left(\ln\frac{v}{v_0}\right)^2\right] . \tag{4.6}$$

The maximum of the distribution is reached at $v_{\max} = v_0\exp(-\sigma^2)$. Expressing this probability distribution in terms of the parameter $v_{\max}$ instead of v_0, we obtain the less familiar but more

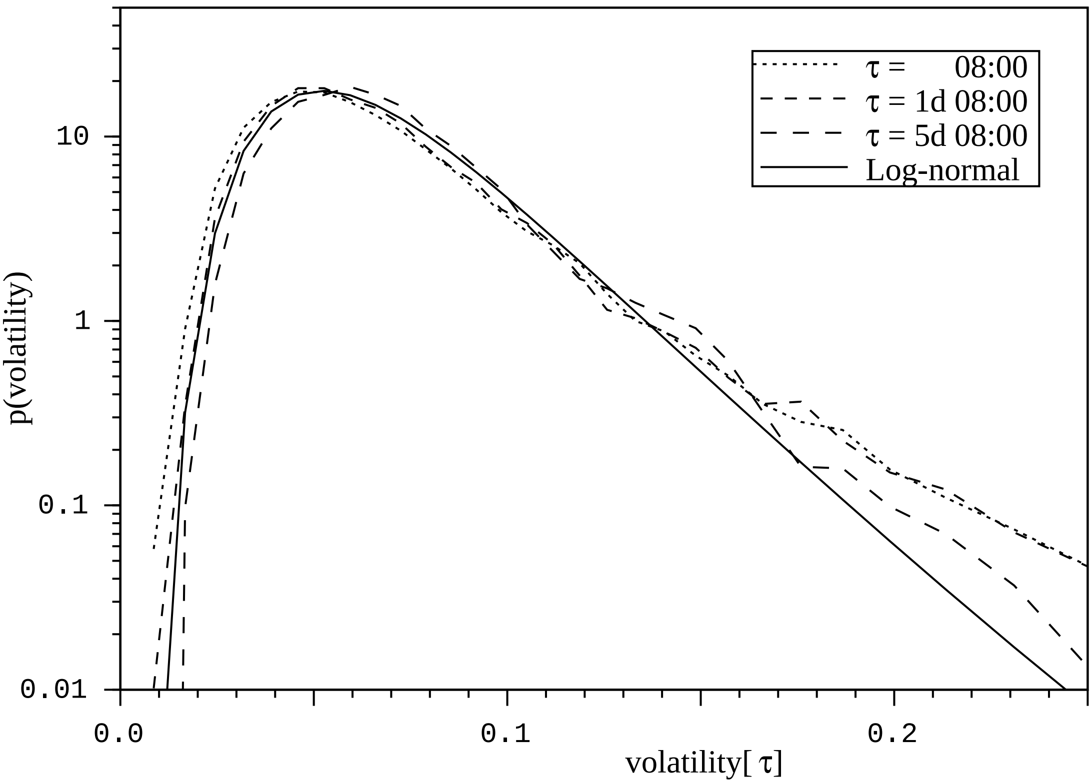

Figure 1: The probability distribution for the volatility. The data is for USD/DEM, from 1.1.1987 to 1.1.1998. The parameters for the fit are $v_m = 0.052$ and $\sigma = 0.4$.

convenient expression

$$p(v) = \frac{1}{\sqrt{2\pi}\sigma v_{\max}} \exp\left[-\frac{1}{2\sigma^2}\left(\ln\frac{v}{v_{\max}}\right)^2 - \frac{\sigma^2}{2}\right] \quad (4.7)$$

$$= p_{\max} \exp\left[-\frac{1}{2\sigma^2}\left(\ln\frac{v}{v_{\max}}\right)^2\right] . \quad (4.8)$$

This form is easier to work with than eq. 4.6. The mean of the log-normal probability density is

$$\overline{v} = v_{\max} \exp(3\sigma^2/2) \ . \quad (4.9)$$

In Figure 1, the measured probability distribution of the volatility is presented, together with a fit to the above theoretical distribution. Clearly, the agreement is excellent for the center of the distributions. It is only in the tail that the observed probability is systematically larger than the theoretical distribution. This is due to the fat tailedness of the return pdf, namely to the fact that $p(r)$ decays as $1/r^{\alpha+1}$ with a tail exponent $\alpha \simeq 3.5$ (Müller et al., 1998). Heuristically, the fat tail of $p(r)$ must translate in a fat tail distribution for $p(v)$ with the same tail exponent. This is not correctly captured by the log-normal pdf. Note also that the Gaussian random walk scaling makes the probability distribution almost invariant when changing the time range τ justifying the volatility definition of eq. 3.5.

We tried to fit the volatility pdf with other 'classical' probability densities (χ, χ^2, F-distribution, Weibull, Fisher-Tippett), but none provided as good a fit as the log-normal. This can be understood partly by the following argument. The log-normal pdf can be rewritten in the form

$$p(v) = p_{\max} \left(\frac{v_{\max}}{v} \right)^{\kappa(v)} \tag{4.10}$$

$$\kappa(v) = \frac{1}{2\sigma^2} \ln \frac{v}{v_{\max}} \tag{4.11}$$

For large v, this law decays faster than any power, but much slower than an exponential or a Gaussian. All the classical pdf's decay too fast, like an exponential or a Gaussian, except the F-distribution which overall does not fit the data well. We can modify the form of $\kappa(v)$ by introducing a saturation at the approximate tail exponent value $\alpha + 1 \simeq 4.5$. We then obtain a good overall fit for the volatility pdf, for example with the form

$$\kappa(v) = \frac{1}{2p\sigma^2} \ln \left(\frac{1}{(v_{\max}/v)^p + \exp(-2p\sigma^2(\alpha+1))} \right) \tag{4.12}$$

with $p = 2$ to $p = 4$. Yet, the cross-over from log-normal to power law behavior seems to be asset dependent. For example, the USD/JPY shows much more tail than DEM/CHF. Without a theoretical hint, such modifications of the theoretical pdf are an *ad hoc* solution, which introduces new parameters. Therefore, we prefer to keep the simpler log-normal form for the volatility distribution.

A puzzling feature of this fit is that the log-normal pdf has two parameters $(v_{\max}, \sigma)$, whereas the return pdf has only one scale parameter measuring the width of the distribution. It appears that we have one additional parameter in $p(v)$. For the sake of argument, let us assume that the return pdf is a Gaussian of width v_0. The standard deviation (volatility) v is defined as $v = \sqrt{1/n \sum r_i^2}$, and is distributed as a rescaled χ pdf $n/v_0\, \chi[n](vn/v_0)$ (a χ distribution because of the square root in the definition of v, and rescaled because of the factor $1/n$ and v_0). This distribution has two parameters n and v_0, yet the parameter n is fixed by the definition of v. Essentially, for the rescaled distribution, the center of the distribution measures v_0 and the width depends on n. The same phenomenon occurs with the log-normal pdf, where v_0 measures the width of the return pdf, and σ depends on the volatility definition. To check for this, we have fitted a log-normal pdf to the empirical distribution of v for many currency pairs, with different $v_{\max}$ but the same $\sigma = 0.4$. All the fits are excellent up to the tails and confirm that there is indeed only one free parameter. Therefore, throughout this paper, we have fixed $\sigma = 0.4$. It is important to understand that since we need to fit many distributions for many different assets, the methodology should be parsimonious and it is a success when we can avoid introducing more parameters.

It is tempting to go one step further and to relate σ to the number of degrees of freedom used in the volatility definition (in our case this number is 16). In Fig. 2, we plotted the χ and the χ^2 distribution with 16 degrees of freedom, the log-normal distribution and the empirical pdf. Unexpectedly, given the above argument, the χ distribution is very far from the empirical one. The χ^2 does better, but the tail is still not fat enough. The best fit is clearly achieved by the log-normal pdf.

A related puzzling feature is the lack of 'central limit theorem'. A one day return is the sum of many returns at a smaller time scale and, if these returns are independent, the condition for the central limit theorem are fulfilled (the tail index must also be larger than 2, which is true at least for FX). The returns are not independent because of the ARCH effect, but empirically, the return

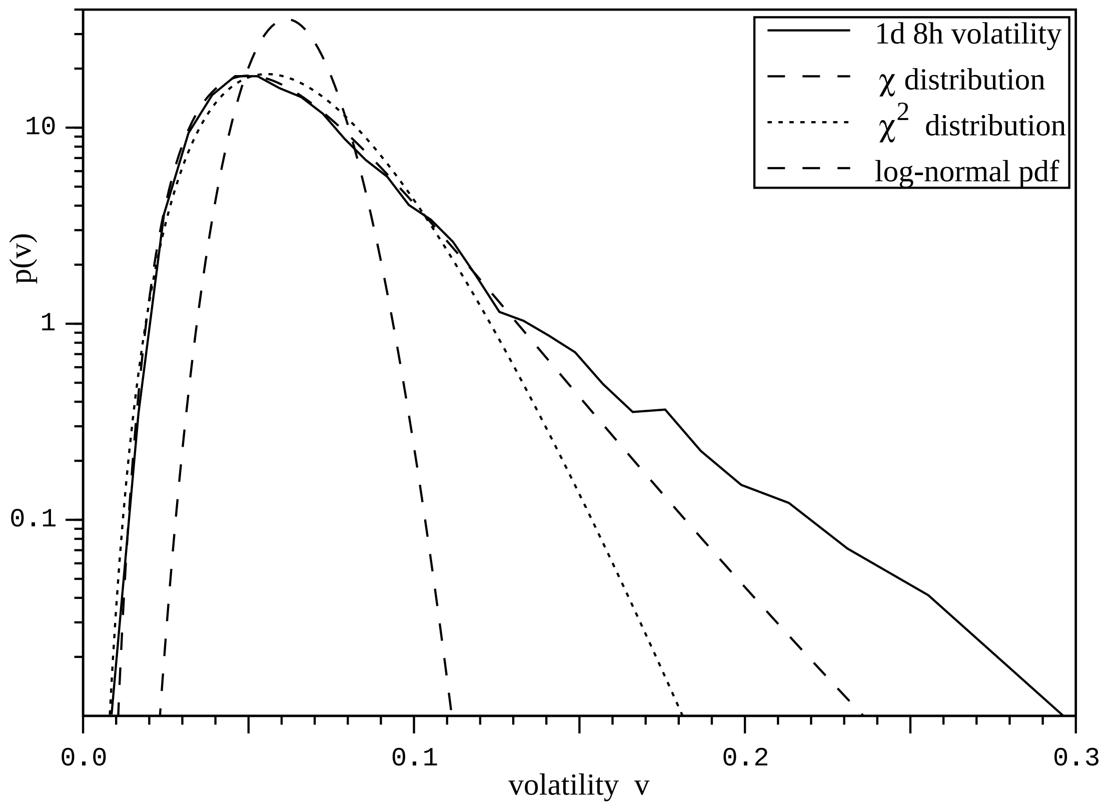

Figure 2: The χ and χ^2 distribution compared to the empirical distribution. The empirical data is for USD/DEM, from 1.1.1987 to 1.1.1998, and with $\tau = $ 1d 8h. The χ and χ^2 distributions are rescaled to have the same mean as the empirical data.

pdf $p(r[\delta t])$ indeed converges towards a Gaussian law with increasing time horizon δt. This fact tells us that the serial dependency in the returns is weak enough for the central limit theorem to apply. Therefore, we might expect the volatility pdf to converge toward a χ distribution with increasing time horizon. Empirically, this is certainly not the case and the empirical volatility distribution instead shows a very remarkable stability with time aggregation. This is probably due to the strong dependency seen in the autocorrelation of absolute returns (Dacorogna et al., 1993; Ding et al., 1993). It is also interesting to relate the volatility distributions with the one of a GARCH(1,1) process. The aggregation law of the GARCH(1,1) process have been derived by (Drost and Nijman, 1993), and essentially the volatility pdf converges toward its mean when aggregating (i.e. the pdf converges toward a delta distribution). This is clearly different from the above findings. Moreover, the other GARCH(1,1) parameters fitted at different time horizons do not seem to follow the aggregation law of the GARCH model (Guillaume et al., 1994; Andersen and Bollerslev, 1997; Zumbach, 1998b). It is a further confirmation that volatility must be measured with different frequencies as argued in the discussion of eq. 2.1.

5 Definition of the Scale of Market Shocks S

In order to fully define the scale of market shocks by the formula 2.1, we must choose the measure $\mu(\ln\tau)$ and the mapping function f.

For the measure $\mu(\ln\tau)$, we take a smooth function, centered at τ_{center}, and decaying toward the short and long time intervals. The exact analytical form does not play an important role. Here, we take

$$\begin{aligned} x_- &= -\alpha_- \ln(\tau/\tau_{\text{center}}) && \text{for} \quad \tau < \tau_{\text{center}} \\ x_+ &= \alpha_+ \ln(\tau/\tau_{\text{center}}) && \text{for} \quad \tau > \tau_{\text{center}} \\ \mu(\ln\tau) &= ce^{-x}\left(1 + x + x^2/2\right) && \text{for} \quad x = x_\pm \end{aligned} \tag{5.13}$$

The parameters α_- and α_+ control the decay of the measure for, respectively, the small and large τ. The constant c is adjusted so that μ is a unit measure $\int d\ln\tau\ \mu(\ln\tau) = 1$. The results presented below are computed with $\alpha_- = \alpha_+ = 2$. Practically, the integral over τ has a cut-off at 1 hour on the low limit and 42 days on the high limit. Since the measure is very small at the integration limit, it does not really influence the result of the integral. The value of τ_{center} is more important and a value around one day gives good results. The parameter τ_{center} is crucial since it controls the time response of the indicator.

In order to construct a good indicator, the form of the function f is very critical. Following a direct analogy with the Richter scale, we first took

$$f_{\ln}(v) = \ln(v/v_{\max}). \tag{5.14}$$

This led to quite a poor indicator, as turbulent markets are not clearly above the normal fluctuating market. In other words, this form for f does not differentiate enough between exceptional events and background fluctuations. A better form for the mapping function can be derived by transforming the usual form of the Richter scale. Empirically, the observed probability of an earthquake with energy E decays as a power law

$$p(E) \simeq \left(\frac{E_0}{E}\right)^{\kappa} \tag{5.15}$$

Using this probability distribution, we can rewrite the Richter scale R as

$$R \simeq \ln\left(\frac{E}{E_0}\right) = \frac{1}{\kappa}\ln\left(\frac{1}{p(E)}\right) \tag{5.16}$$

In this form, the Richter scale is expressed as the logarithm of the (un)likelihood of an event with energy E. This suggests the following form for the mapping f

$$\begin{aligned} f_{\text{adp}}(v) &= \text{sign}(v - v_{\max}) \ln\left(\frac{p_{\max}}{p(v)}\right) \\ &= \text{sign}(v - v_{\max}) \frac{1}{2\sigma^2}\left(\ln \frac{v}{v_{\max}}\right)^2 \end{aligned} \tag{5.17}$$

where, in the last equality, we use the fit of the previous section for the probability distribution $p(v)$. These two definitions for f eq. 5.14 and eq. 5.17 lead to different indicators, because for financial data the probability distribution of the volatility v is not a power law, contrary to the underlying process of the Richter scale. The theoretical probability distribution for the volatility and both mappings f are presented in Figure 3. We clearly see why the second definition is working well: the range of normal volatilities around the maximum $v_{\max}$ are mapped closely to zero, and therefore are scaled down. On the other hand, the mapping is very linear in the large volatility region, extracting the important events from the noise. The mapping 5.17 acts as a non linear amplifier that extracts the signal we want from the normal volatility level. This construction is efficient because financial data are dominated by the white noise component. In this respect, finance is very different from geophysics, where seismograph recordings show with great clarity the passage of the various sismic waves.

In eq. 5.17, we use the fit for the volatility pdf. Improving on this fit could lead to a better mapping. As argued before, we prefer to keep a simple log-normal pdf in order to avoid introducing new asset dependent parameters in this fit. Compared to empirical data, the log-normal pdf underestimates the tail probability. In turn, this gives too much weight to large events compared to $f(v) \sim \ln p_{\max}/p(v)$. As we are precisely interested in extreme events, it is indeed a good feature to slightly overemphasize large volatilities.

As argued above, the parameter σ is fixed by our volatility definition and therefore does not depend on the asset. Yet, the parameters $v_{\max}$ introduces an explicit dependency on the asset in the definition of S. Here we reach a difference between geophysics and finance: there is only one earth and one Richter scale whereas different assets are like different planets. It is intuitively clear that the USD/DEM experiences more fluctuations and shocks than DEM/CHF, and the difference is even larger, say, between FX and interest rates. At this point, two approaches are possible:

- We want to compare the scale of market shocks values for various assets. For this purpose, a universal SMS is needed.
- We are interested in individual assets and want a scale calibrated to each asset behavior. For this purpose, an adaptive SMS is needed.

As both approaches have their respective merits, let us introduce two scales of market shocks: the universal SMS S_{uni} and the adaptive SMS S_{adp}.

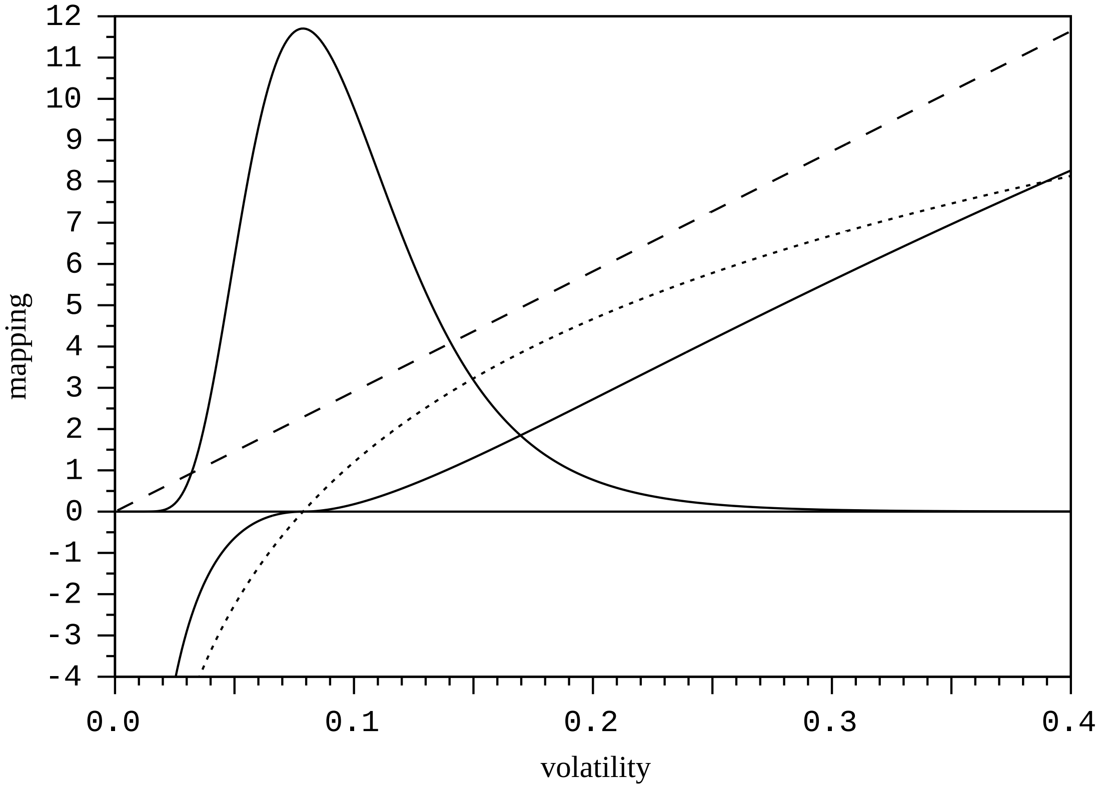

Figure 3: Different choices of mapping functions f and the fitted theoretical probability distribution for the volatility are shown. We plot the mappings f_{ln} (5.14) (dotted line), f_{adp} (5.17) (full line) and f_{uni} (5.18) (dashed line). For presentation purpose, the mapping (5.14) has been multiplied by 5. The parameters are $\overline{v} = 0.1$ and $\sigma = 0.4$. With these parameters, the slope of the mappings f_{adp} and f_{uni} are equal.

- **The universal SMS S_{uni}**
  The mapping (5.17) is almost linear for the large volatilities in which we are mainly interested. Therefore, a simple linear mapping already provides a good scale of market shocks

$$f_{\text{uni}}(v) = sv. \tag{5.18}$$

  The slope s is chosen so that an asset with a 10% annual volatility has comparable SMS differential values in both scales at $v = 3v_{\max}$

$$s = \left.\frac{df_{\text{adp}}(v)}{dv}\right|_{v=3v_{\max}} = \frac{\ln(3)}{\sigma^2 3\, v_{\max}} = 29.1 \tag{5.19}$$

  with $\overline{v} = 0.1$, $v_{\max} = \overline{v}\exp(-3\sigma^2/2) =$, and $f(v)$ given by eq. 5.17. With this mapping, the background normal volatilities $\overline{v}$ give a value $S_{\text{uni}} = s\overline{v}$ depending on the asset.

- **The adaptive SMS S_{adp}**
  This scale is defined with the mapping (5.17). Yet, for each asset, this mapping depends on the mean volatility $\overline{v}$ which has to be estimated. Moreover, this quantity is not constant over a time scale of a few years. For example, the intra European FX market has a decreasing volatility since the beginning of the '70s and a large event in 1998 for DEM/FRF seen in the adaptive SMS is certainly much smaller in the absolute SMS. Therefore, the mean volatility $\overline{v}$ is like an adiabatic, slowly changing quantity, namely it is a time series. We measure $\overline{v}$ with

$$\overline{v} = \int d\ln\tau\; \mu(\ln\tau)\; \text{MNorm}[90d, p = 1/2; v[\tau; x]] \tag{5.20}$$

  The τ integration and the measure are as for the definition of the SMS. The MNorm operator (Zumbach, 1998a) computes a moving norm with $\tau = 90$ days, corresponding approximatively to a (rectangular) moving window of length 180 days. The norm is computed with an exponent $p = 1/2$ in order to decrease the importance of the large volatilities and to obtain a more robust estimate of the mean volatility. With this definition for $\overline{v}$ and the relation 4.9, we now have a complete definition for the adaptive SMS S_{adp}.

  The possible gaps in the data source lead to a fictitious small value for the volatility. Because of the logarithm in the mapping f, data gaps would produce negative diverging values for f, given in turn a strongly negative SMS indicator S. In order to avoid this problem being due to the data source, the mapping is bounded below, namely if $v < v_{\min}$ then $f_{\text{adp}}(v) = f_{\text{adp}}(v_{\min})$. In the present computations, we take $v_{\min} = \overline{v}/3$.

For the software implementation, the value of the volatility $v[\tau]$ is updated with every tick of the market. Then, after every time interval ΔT, the index S is computed. The integral over τ is computed on a regular logarithmic grid, ranging from 1 hour to 42 days, with a reason of $2^{1/4}$. The statistical properties of S can then be studied from this regular time series. Furthermore, in order to compute the probability distribution for the volatilities and for S, the weekends have been removed.

6 A first empirical study

In order to show the behavior of both scales of market shocks, we display in Figure 4 the indices S_{uni} and S_{adp} for the USD/DEM and for DEM/CHF. These particular currencies pairs are

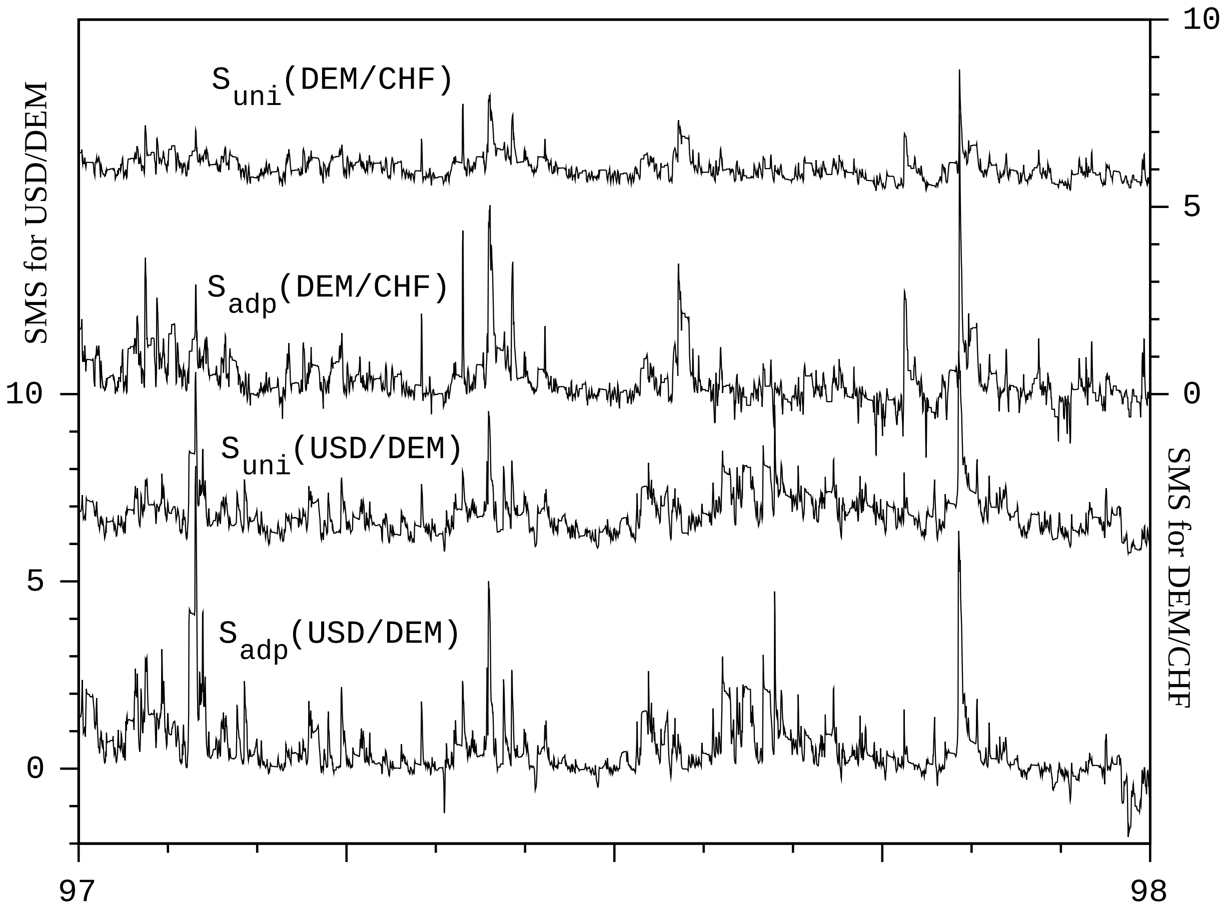

Figure 4: The SMS for USD/DEM (left scale) and for DEM/CHF (right scale) for the year 1997. Both universal scales have been shifted by 5 points for presentation purpose. The horizontal axis is divided into 12 equal intervals, approximating one month.

selected because the CHF is strongly correlated with the DEM, therefore a very low level of volatility is expected. Both adaptive scales have zero as their base lines, and show comparable peak heights. Clearly, the adaptive mechanism put into $f_{\mathrm{adp}}(v)$ is working. Yet, the overall lower volatility of DEM/CHF is clearly visible on S_{uni}, with a base line of ~ 1, whereas the base line for USD/DEM is ~ 2. The peak heights is also much smaller for DEM/CHF, in line with our expectation for a direct comparison of these currencies pairs.

As another illustration of the SMS indicator, we compute its value for the USD/JPY market, from January 1997 to October 1998. This is displayed in Figure 5. We choose this example because 1997 and 1998 were eventful years for the JPY and it is recent enough so that the reader may remember the events causing the peaks (hint: some events are due to other currencies, like the Thai Baht). For the adaptive SMS S_{adp}, the non-linear transformation f is working very well, as peaks distinguish themselves clearly from the background fluctuations. It should also be emphasized that the volatilities used to compute S are mostly not visible on the price curve. The volatilities with the largest weight corresponds to intra-day fluctuations and on this graph there are only 6 points per day $\Delta T = 4$ hours. Therefore, the scale of market shocks really brings out new information, not visible in the price curve.

We also study the dependency of S on the center of the measure τ_{center}. The main conclusion is that varying τ_{center} from 12 hours to 48 hours does not significantly change the results. As

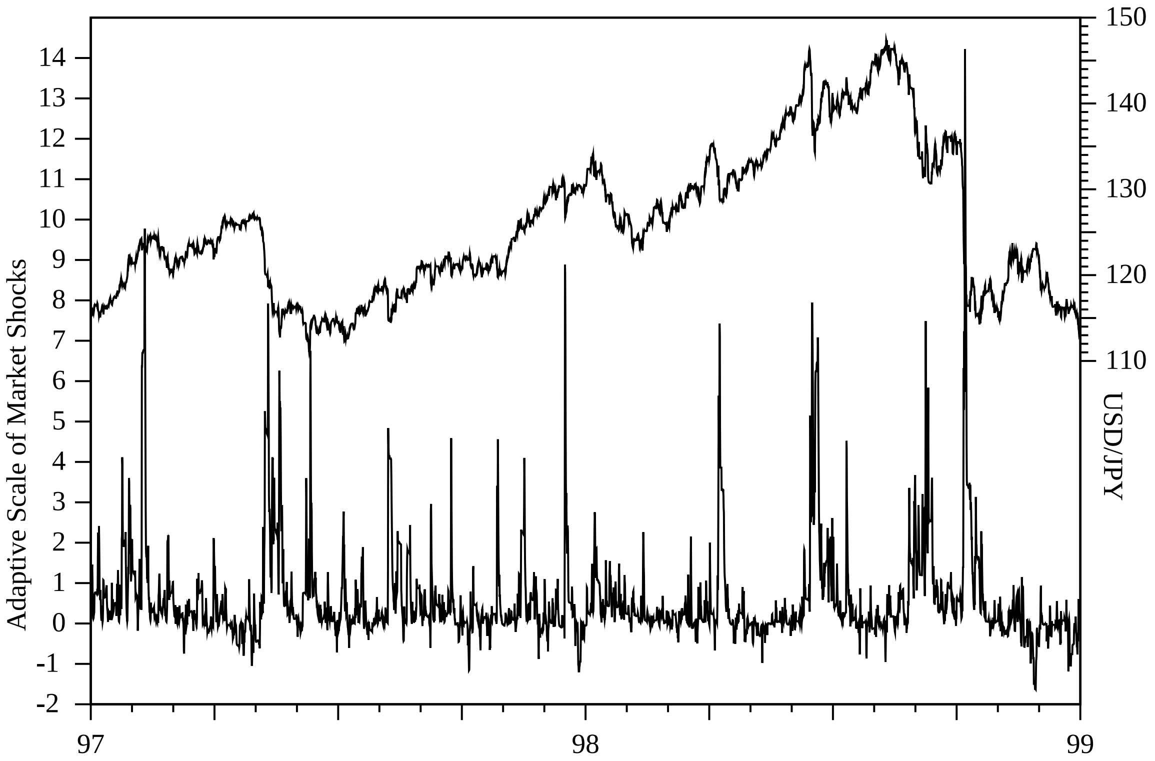

Figure 5: The Scale of Market Shocks for USD/JPY (left scale) and the corresponding price (right scale). The horizontal axis is divided into 12 equal intervals, approximating one month.

can be expected, by taking a smaller τ_{center}, the peaks tend to be slightly narrower and higher. In the previous graph, the only exception are the high narrow peak in February and the two small narrow peaks in July 1997: those peaks appear to be dominated by very short intra day fluctuations.

The empirical pdf's for both indices can be measured for different currency pairs. The result for USD/DEM is presented in Fig. 6, as well as a reasonable fit for both pdf(S_{uni}) and pdf(S_{adp}) in the region of large S. The functional form of the fits is not necessarily the best for this particular currency pair, but provides a reasonable fit for all currency pairs. The asymptotic properties of the pdf for the universal index is well described by a power law decay, with an exponent varying from 5.5 to 7. This is consistent with the functional form of the definition and with the asymptotic decay for the return pdf, but with slightly larger tail exponents. Yet, the value of the decay exponent should be taken with care when comparing to a tail exponent, because we only applied a fit for the large S region and not a proper tail estimate. The pdf for S_{adp} is well approximated by the exponential form $\exp(-\alpha S_{\text{adp}})$, with a coefficient α between 0.55 to 0.75. This is roughly consistent with the rule of thumb that one more point on the (adaptive) scale of market shocks corresponds to an event that is twice as unlikely. Yet the precise values are asset dependent. This result justifies *a posteriori* the argument and mapping used to construct the adaptive scale.

7 Event study for two FX-rate scales

To illustrate the quality of the scale, we choose to compute it for USD/JPY and USD/DEM from January 1997 to September 1998. These particular months were full of events in East

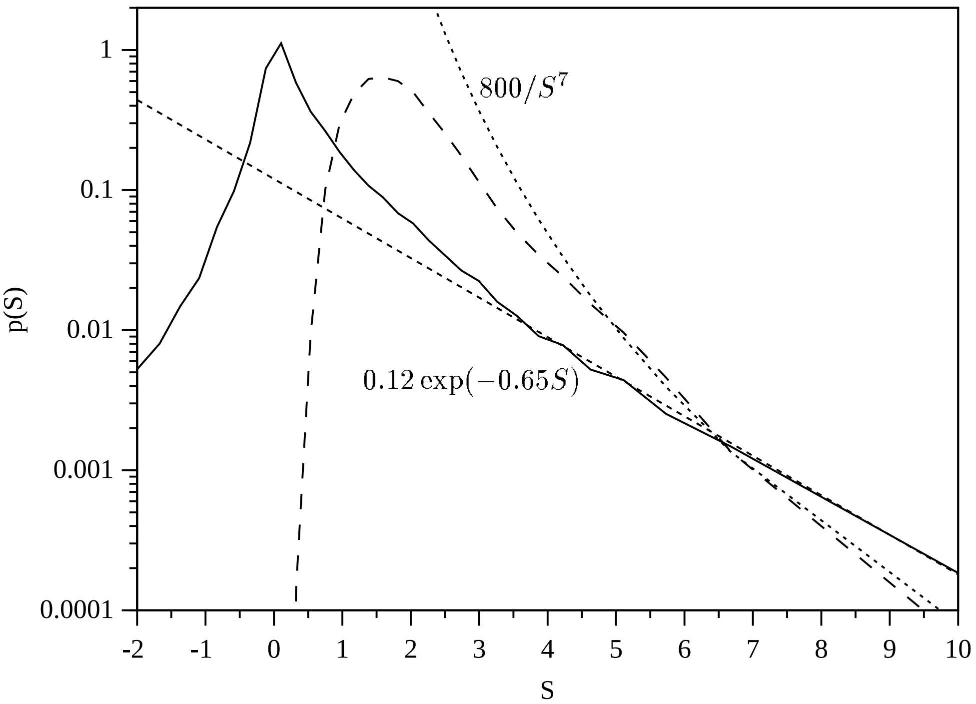

Figure 6: Probability distributions for S_{uni} (dashed line) and S_{adp} (full line), with the asymptotic fits (dotted line). The data is for USD/DEM, from 01.01.1988 to 01.12.1998. The pdf are constructed with data sampled every 4h in ϑ-time.

Asia and, by contrast, relatively calm in Europe. Thus comparing the two behaviors should give us an idea whether the reality is well described. In Tables 1 and 2, we report the days where the universal SMS for USD/JPY or for USD/DEM were above 3.0. There are close to 40 of these days for USD/JPY over 19 periods (some consecutive days) while there are only 17 over only 8 periods for USD/DEM. To find the events corresponding to the peaks, we looked among the money market news headlines of Reuters (AAMM page) and in the magazine "The Economist" for reports of particular news during those days. The Reuters AAMM page was discontinued after September 1998, forbiding to extend this study further. From the table, it is clear that the rumors of the intervention of the Bank of Japan, which we found publicized in Reuters headlines, had a significant impact on the scale together with the major events that hit the Japanese market, which were reported both in "The Economist" and on Reuters. There is one noticeable exception: the resignation of the governor and the deputy governor of the bank of Japan (only 2.2 on the scale). The reason for this relatively small movement is probably the fact that these resignations were largely expected after the turmoils of the previous year. The highest value (9.9) was reached on December 17, 1997 after the announcement of the government plan to restore stability. This plan was a big disappointment for investors and the sentiment over the authorities' failure to get out of the financial crisis became very strong. In contrast, the JPY was comparatively less affected by the October crisis on the Asian stock markets.

Similarly, the few days where the scale was high for the USD/DEM are, as expected, related to the stock exchanges turbulences after the Hang Seng crash at the end of October 1997 and also in August 1998. In October 1997 the Hang Seng crash also affected the European and American exchanges (the DAX was down 6.6% on October 28). The scale reaches its peak (6.1) for 1997 on October 28 during the Asian exchange crisis. The SMS for USD/DEM is also quite sensitive to the Bundesbank changing or keeping the REPO rate (second largest SMS 5.6 in February 1997 when the Bundesbank set the REPO rate to 3%).

Date	JPY	DEM	News
08.02.97	5.9	3.5	Article in the Economist from Feb. 15-21 (p. 79): problems with Nippon Credit and Hokkaido Takushoku bank. "Nippon Credit admits 11.4 billions USD in bad loans analysts think it is the double".
09.02.97	5.9	3.4	
10.02.97	9.1	5.6	Rumors of BOJ (Bank of Japan) intervention for 500 bln JPY. The Bundesbank sets the REPO rate at 3%.
09.05.97	5.1		Thai Baht crisis. Saved from devaluation by central
10.05.97	4.9		bank interventions.
11.05.97	4.8		
12.05.97	7.1		Rumors of BOJ intervention for 400 bln JPY.
.			
.			
21.05.97	5.8	4.5	Comments of BOJ deputy governor in parliament.
09.06.97	4.8		Rumors of BOJ intervention for 100 bln JPY. Bank of Korea intervenes in interbank market, buys.
.			
11.06.97			Rumors of BOJ (Bank of Japan) intervention for 300 bln JPY.
12.06.97	8.0		
13.06.97	4.6		Japan prosecutors arrest DKB (Daichi Kangyo) vice president.
08.08.97	6.1		Taiwan dollar sinks to 28.739 at close.
09.08.97	5.4		
10.08.97	5.4		
11.08.97	5.3		
26.08.97		4.7	The Bundesbank leaves the REPO rate unchanged contrary to market expectations.

Table 1: Table summarizing the dates where the universal SMS for either USD/JPY or USD/DEM were above 3 and the related events, as found in "The Economist" and in the monetary news headlines of Reuters. The second and the third column report the highest value reached by the SMS during that day for USD/JPY and USD/DEM respectively. An empty entry means that the scale was below 3.

There is a substancial body of litterature on the relations between news and market movements. Most of these papers use low frequency data (daily, sometime hourly), look at price differences (return) and focus on economic news. The major difficulty is to obtain an *a priori* quantitative estimate of the announcements, possibly separating the expected and unexpected part of announcements. Generally, the relation is found to be very small or non existent. In a recent paper, Almeida, Goodhart and Payne (Almeida et al., 1998) studied the correlation between high frequency price movements and the unexpected part of macroeconomic announcements for USD/DEM. They find a small (at most 30 basis points) but significant impact for most announcements, on very short time horizon (15 minutes). The correlations decay rapidly with increasing time horizons, and is unobservable at one day. In the conclusion, they conjecture that from the macroeconomic news, the unexpected changes in interest rates should produce the largest responses on the exchange rates [1]. This is indeed what we observe in Tables 1 and 2. Compared to the bulk of the litterature, the present approach is different as we focus on high frequency volatility at different time horizons. This allows us to obtain clear signals and a convincing relation between large values for the SMS indexes and major events. Yet, a quantitative study of the correlation with the unexpected part of the news remains to be done. The major obstacle is in quantifying the expected and unexpected part of news for a broad spectrum of events.

[1] We discovered this paper after our event study was completed.

Date	JPY	DEM	News
09.09.97	4.9		USD hit by fears of US/Japan trade tensions.
24.09.97	6.3		Resignation of the chairman of Daiwa Securities. Arrest of a former president of Yamaichi.
28.10.97	6.0	6.1	Stock crash originating in Asia (HSI -13.7%, DJI -7.2%, DAX -6.6%).
29.10.97	4.5	4.8	
17.11.97	5.5		Collapse of Hokkaido Takushoku Bank (10th largest com. bank).
17.12.97	9.9		Announcement of the government "measures to restore the path of stability". The investors are disappointed by the low tax cut announced. The BOJ is said to intervene with more than 1 bln USD.
18.12.97	4.8		The BOJ is believed to be selling USD at 127.50/60. The grain brokerage firm Toshoku files for bankruptcy.
08.01.98	3.6		Hong Kong Peregrine's bankruptcy. Indonesian Rupiah Crisis. BOJ intervention after JPY decline.
26.01.98		4.5	Possible Clinton's affair with a former White House employee is made public on internet.
20.03.98	2.2		Resignation of the head of Bank of Japan (Yasuo Matsushita) and of his deputy (Toshiko Fukui)
09.04.98	7.5		BOJ announces intervention: selling USD in New York.
10.04.98	9.1		Analysts think that BOJ spent 5 bln USD in the last two days. BOJ disciplines 98 staff members for "entertainment". BOJ officials admit leaking internal information to contacts.
13.04.98	5.2		Announcement of Japan trade surplus that jumps up 97.1%.
10.05.98		5.1	The SPD admits that it could rule with communists
11.05.98		5.5	India tests an atomic bomb
16.06.98	6.2		
17.06.98	9.0		Joint intervention of Bank of Japan and US Federal Reserve
18.06.98	7.8		
19.06.98	6.2		
20.06.98	7.4		G7 Meeting on Asia Crisis
21.06.98	7.6		
22.06.98	8.6		USD again up on disappointment from the G7 meeting
13.07.98	6.7		Hashimoto's defeat at the election
28.08.98	5.6	3.7	Turmoil on the stock markets
29.08.98		3.6	
30.08.98		3.5	
31.08.98		4.4	
01.09.98	6.4	4.2	
02.09.98	5.3	3.3	
04.09.98	5.7	3.4	

Table 2: Table summarizing the dates where the universal SMS for either USD/JPY or USD/DEM were above 3 and the related events, as found in "The Economist" and in the monetary news headlines of Reuters (continuation from Table 1). The second and the third column report the highest value reached by the SMS during that day for USD/JPY and USD/DEM respectively. An empty entry means that the scale was below 3.

8 From a FX-rate scale to a 'Grand' Market Scale

The scale of market shocks can be constructed in principle for any market: the index is computed from the price time series. In the foreign exchange (FX) market, an index S[per/exchanged] is associated to each currency pair. Yet, in the case of the FX market, it is interesting to derive an index per currency. For example, when S[USD/JPY] is large, you cannot determine if the turbulence is originating in the US or in Japan. By considering more currency pairs, like USD/DEM, USD/GPB, USD/JPY, etc..., the USD part can be isolated. By summing currency pairs, the contributions of one country is enhanced while reducing the effect of the other currencies. A single currency index S[per] for the 'per' currency is computed by summing the S[per/exchanged] index for currency pairs 'per/exchanged' with respect to the 'exchanged' currency. Each currency pair in the sum is weighted by its estimated relative importance.

This procedure is illustrated in Figure 7 for the DEM and USD, with a small currency basket. The weights for the single currency SMS are

$$\begin{aligned} \text{USD} &= 0.3\ \text{USD/DEM} + 0.3\ \text{USD/JPY} + 0.2\ \text{GBP/USD} + 0.1\ \text{USD/CHF} \\ \text{DEM} &= 0.4\ \text{USD/DEM} + 0.3\ \text{DEM/JPY} + 0.2\ \text{GBP/DEM} + 0.1\ \text{DEM/CHF} \end{aligned}$$

where we have used the symmetry in the 'per' and 'exchange' currencies. In order to isolate better the contribution of one currency, the currency pairs basket should be as large as possible. For example, for the USD, contribution from other asiatic currencies and from central and south-american currencies (Mexican pesos, Brasilian Real) can also be included.

Finally, the single currency indices S can be summed to obtain a world index, reflecting the total currency turbulences in the worldwide FX markets. Each contributing currency is weighted according to its estimated relative importance.

9 Forecasting Price Movements

Let us emphasize that both scales of market shocks are designed as indicators: they provide a diagnostic tool about the past behavior of the price movements. They can be seen as measuring a 'state' of the market and they allow to compare states for different times and assets. In order to have a fast enough response, the integration measure is centered at one day volatility, corresponding to a sum of squared returns measured on price movements over 1.5 hours. These short time intervals result in the indicators having a quick intra day response, and an accurate localization in time of the events. We find a good illustration of this behavior in the exceptionally strong price movement of USD/JPY during the week of October 5-9, 1998. The movement was of more than 15% in one week, the largest weekly movement recorded in our database, and happened in two consecutive shocks, one on October 7th and one on October 8th. The SMS moved from a value of below 3 to a value above 10 in a few hours before the movement was completed thus giving an early warning that the situation was very unstable.

The scale of market shocks has not been designed as a forecast of volatility. Moreover, due to the strong relation between shocks and external events, a poor forecastability can be expected *a priori*. Nevertheless, it is interesting to measure the relation between the SMS and the size of the next price movement to assess its capacity to detect possible risky situations. Therefore, we measure the linear correlation coefficient between the universal SMS and the absolute value of

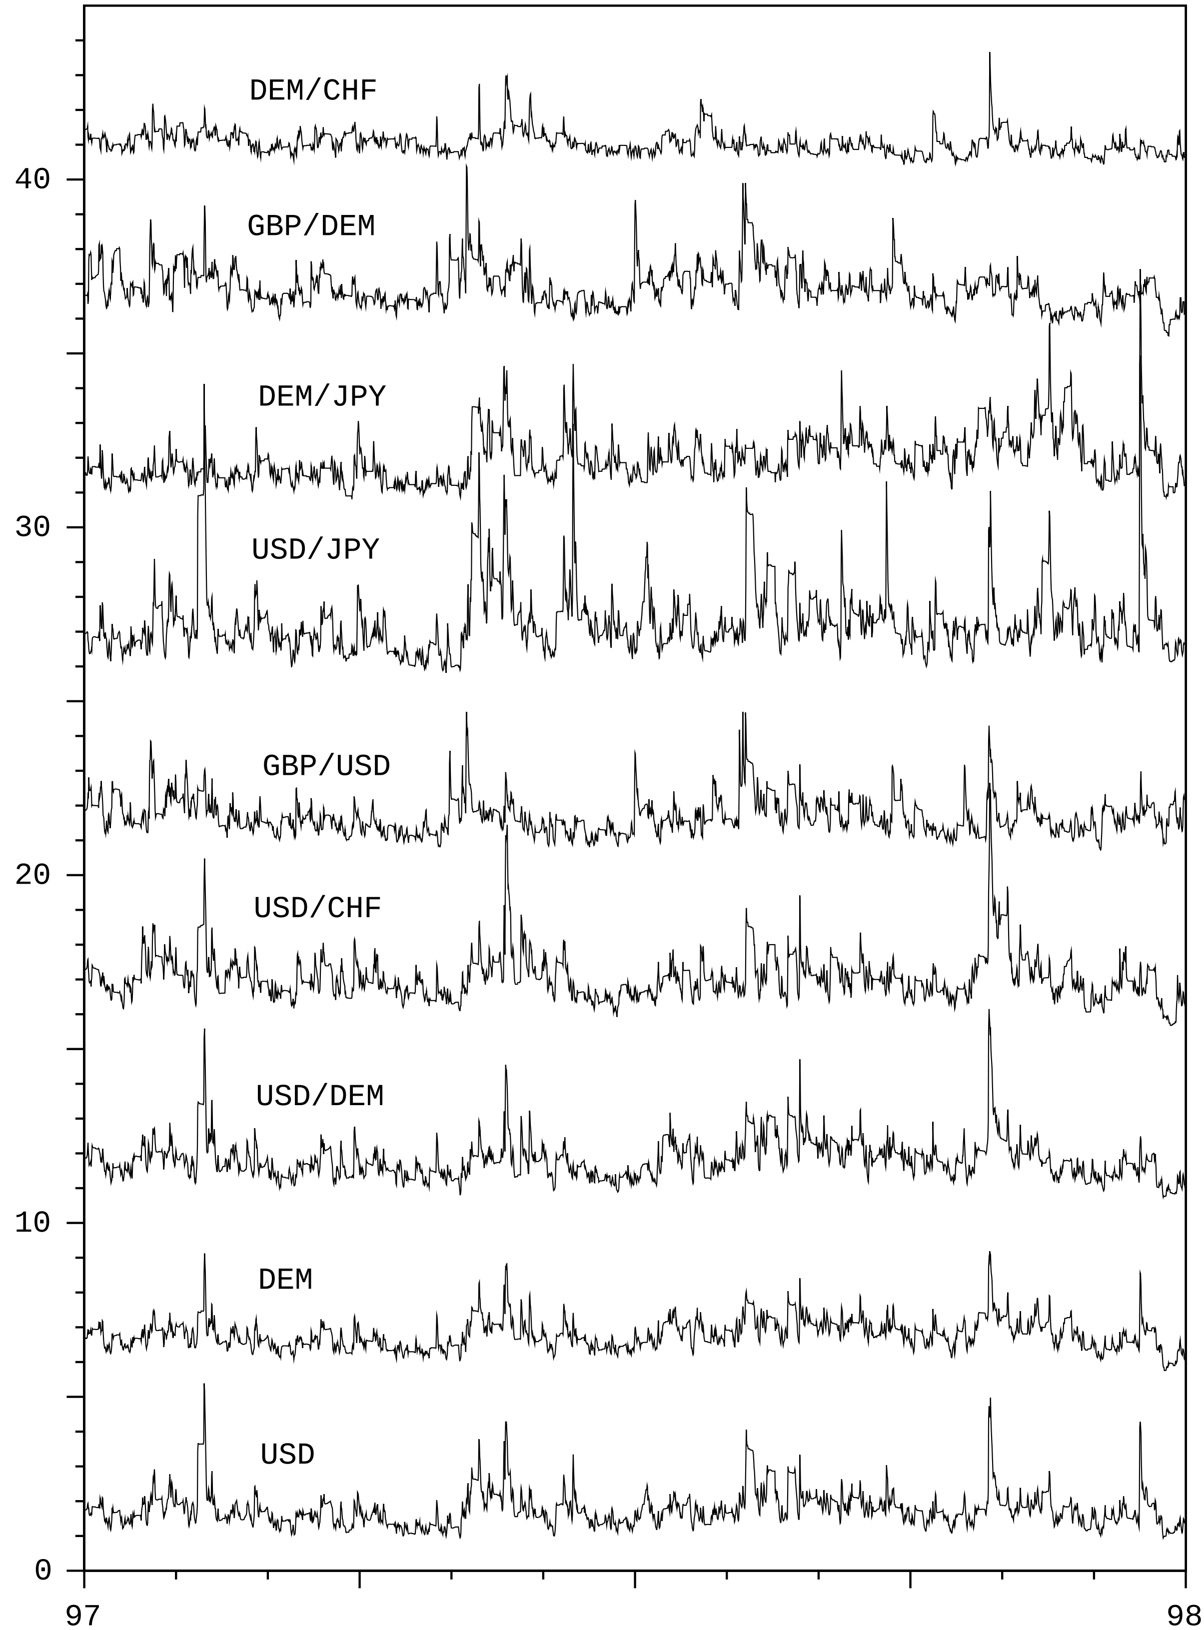

Figure 7: The universal Scale of Market Shocks S_{uni} for the currency pairs DEM/CHF, GBP/DEM, DEM/JPY, USD/JPY, GBP/USD, USD/CHF, USD/DEM and for the DEM and USD. For presentation purpose, each curve is shifted by 5 points. The horizontal axis is divided into 12 equal intervals, approximating one month.

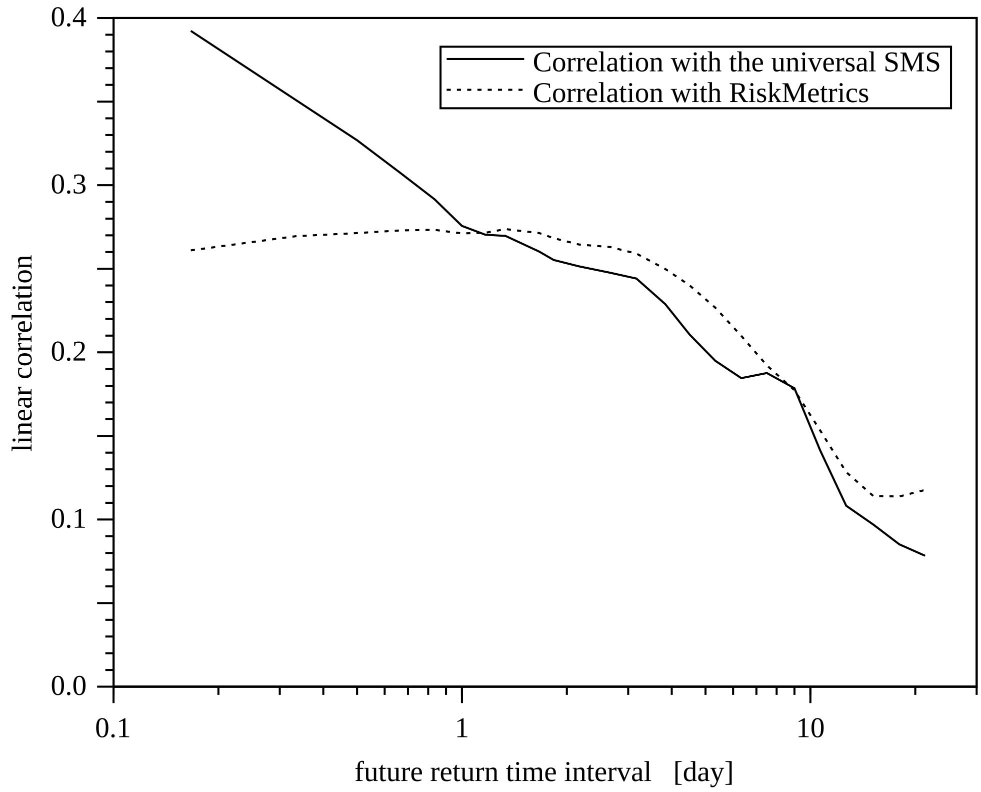

Figure 8: The linear correlation of the next return at time scale τ (horizontal axis) with the Scale of Market Shocks and with a RiskMetrics-like volatility. The data is the FX rate series for USD/JPY, from 01.01.1988 to 01.12.1998. The SMS and RiskMetrics volatility are computed with high frequency data, then sampled every 4h in theta-time. The next return and correlations are computed over the 4h sampled data.

the next return. Essentially, this is similar to a correlation between past and future volatilities. In order to compare the result to a reference curve, we do a similar computation with a volatility model *à la* RiskMetrics (Morgan Guaranty, 1996). In order to obtain a high frequency volatility similar to RiskMetrics, we compute

$$v_{\mathrm{RiskMetrics}} = \sqrt{\mathrm{EMA}[T;(D[\tau;x])^2]}. \tag{9.21}$$

with $\tau = 7/5$ days, $T = 7/5{*}16.16$ days, and 16.16 days $= -1/ln(0.94)$. The operators are described in detail in (Zumbach, 1998a). The computations are done in theta time to remove the seasonalities, and 7/5 is the conversion factor of one physical day into one business day. The original RiskMetrics definition is based on daily prices, our definition uses high-frequency data. We thus expect our measure of volatility to be slightly better than the original RiskMetrics measure. Note that the RiskMetrics parameter 0.94 was optimized to provide the best one day volatility forecast.

The linear correlation coefficient is given in Fig. 8 for both models. Clearly, the SMS is a good leading indicator for the size of the next price movements, particularly for the very short time horizons where the correlation coefficient reaches 0.4 for 4 hours movements. The SMS has a larger correlation than RiskMetrics (by about 13% for 4 hours) up to one day, and a slightly lower correlation for longer return time intervals. This is to be expected given the different time characteristics of the two indicators. A rapid test for a SMS centered at roughly the same time horizon as RiskMetrics (16 days) shows a better correlation for the SMS than for the RiskMetrics volatility. This confirms that our formalism is able to catch risky situations and the return heteroskedasticity, although its main purpose is to be able to quantify and compare *current* price movements. More research in this area is needed to assess to what extent large movements can be expected when the SMS is already indicating values higher than usual.

10 Conclusion

By analogy with the Richter scale in geology, we suggest defining two scales of market shocks based on volatilities of financial asset measured at different frequencies. A careful design of these scales enable us to extract from the price time series the major events hapening on the markets. In a simple event study for USD/JPY and USD/DEM over the last 21 months, we observe that the SMS indexes measure well the market evolution, breaks or crises in Europe, USA and Asia during this period. Therefore, the SMS indices allow us to measure in an objective way the relative impact of different news and events on the overall foreign exchange market. Moreover, the scales for all FX-rates related to a given currency can be combined into a single currency scale which, in turn, can be combined into a general world scale for the FX market.

The definition of the SMS is a mandatory step in evaluating the *current* state of financial markets, and it allows us to quantify in an objective way events and crises. The SMS measured across multiple assets, and some natural combination of them, would let us assess the extent and severity of a crisis. The next step would be to construct a forecast of this indicator, or of other similar quantities like the size of the following price movements. An analysis of the correlation coefficient between the universal SMS and the *following* absolute return shows that the SMS can already be used as an indication of market instabilities in the near future. The ultimate goal would be to have a better forecast of possible future crises, and therefore provide an early warning about potential turmoil on financial markets.

Acknowledgments: The authors would like to thank Leena Tirkkonen and Devon Bowen for providing assistance with the software, Jonathan Buchanan for a careful review of the manuscript and the research group at O&A for interesting and lively discussions. GOZ thanks Pierre Ihmle for sharing his knowledge about earthquakes.

References

Almeida A., Goodhart C., and Payne R., 1998, *The effects of macroeconomic news on high frequency exchange rate behavior*, Journal of Financial and Quantitative Analysis, **33**(3), 383–408.

Andersen T. G. and Bollerslev T., 1997, *Intraday periodicity and volatility persistence in financial markets*, Journal of Empirical Finance, **4**(2-3), 115–158.

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

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.

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.

Guillaume D. M., Dacorogna M. M., and Pictet O. V., 1994, *On the intra-daily performance of garch processes*, Poster presented at the First International Conference on High Frequency Data in Finance, Zürich.

Hsieh D. A., 1988, *The statistical properties of daily foreign exchange rates: 1974-1983*, Journal of International Economics, **24**, 129–145.

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

Morgan Guaranty, 1996, *RiskMetrics^{TM} – Technical Document*, Morgan Guaranty Trust Company of New York, New York, 4th edition.

Müller U. A., Dacorogna M. M., Davé R. D., Olsen R. B., Pictet O. V., and von Weizsäcker J. E., 1997, *Volatilities of different time resolutions – analyzing the dynamics of market components*, Journal of Empirical Finance, **4**(2-3), 213–239.

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.

Müller U. A., Dacorogna M. M., and Pictet O. V., 1998, *Heavy tails in high-frequency financial data*, published in the book "A practical guide to heavy tails:"Statistical Techniques for Analysing Heavy Tailed Distributions", edited by Robert J. Adler, Raisa E. Feldman and Murad S. Taqqu and published by Birkhäuser, Boston 1998.

Richter C. F., 1958, *Elementary Seismology*, Freeman, San Francisco.

Taylor S. J., 1986, *Modelling Financial Time Series*, J. Wiley & Sons, Chichester.

Zumbach G., 1998a, *Operators on inhomogeneous time series*, Internal document GOZ.1998-11-01, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

Zumbach G., 1998b, *The pitfalls in fitting garch processes*, Internal document GOZ.1998-05-07, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.