

Shocks in the S&P500*

Remco T. Peters[†]

December 11, 2003

Abstract

It is well known that financial markets exhibit periods of high activity and low activity. In the present paper high frequency data of the S&P 500 futures stock index from January, 1988 until August 2001 is used to analyze this phenomenon. The aim is to describe the activity in terms of piecewise linear functions, and to relate news items to transition epochs.

*R.P. gratefully acknowledges the financial support of AOT N.V. The author is indebted to Guus Balkema and Robin de Vilder for their detailed comments and suggestions.

[†]email: remco@science.uva.nl. Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Plantage Muidergracht 24 1018TV Amsterdam, The Netherlands. Tel. +31 20 5256521.

1 Introduction

High frequency data for financial processes, in the setting of diffusion models, allow one to obtain a good estimate of the quadratic variation of the process and hence of the volatility as a function of time. In the classical Black and Scholes set-up the squared volatility is the derivative of the quadratic variation of the log price process. This relation remains valid for general diffusion models.

For the S&P 500 futures index, over the period 1988-2001, using 650 000 observations, one obtains for the quadratic variation the continuous, increasing function \hat{q} in figure 1. In this graph one may distinguish regimes during which the slope is approximately constant. At certain points, marked b for breakpoint in the graph, we see abrupt changes in the slope of the function. At other points on the graph there is a change in the slope which is less abrupt. Occasionally there is a short period of very high activity, marked by \mathbb{b} .

The graph of \hat{q} in figure 1 was obtained in Peters and de Vilder (2002a). It raises two questions.

1) Is it possible to use statistics in order to quantify the visual information present in the graph of \hat{q} ? There is a large literature on estimating the location of a single breakpoint, but for an indefinite number of breakpoints less is known. The method used in this paper is inspired by Bai and Perron (2001) and based on Fisher (1958). Very little is known about the volatility process. Indeed it is only with the advent of high frequency data that one can observe this process. Our statistical analysis is based on the assumption that the observed sample function in figure 1 is the sum of a piecewise linear function and a noise process δB where B is a standard Brownian motion and δ a positive constant. There are indications that δ may vary from regime to regime but in this paper we have opted for a simple model, and taken δ constant.

2) If the statistical analysis supports the visual impression of a piecewise linear function then one should be able to give confidence intervals for the breakpoints, thus creating the possibility of linking these points to real world events.

It remains unclear how one should interpret a piecewise constant volatility in terms of economic theory. An analogy which comes to mind is quantum physics where the state of the system is a discrete variable and the system may jump from one state to a higher or lower state by absorbing or emitting

a photon. A look at figure 13 will help to clarify this analogy.

We shall discuss the relevant literature briefly now. Hamilton (1989) focuses on estimating regime switching models with a finite number of states. See also Franses and van Dijk (2000) and their references on regime switching models and their application to volatility. Several authors have commented on the presence of breakpoints (also called change points) in financial data. Timmerman (2001) discusses the existence of breakpoints in the dividend process of U.S. stocks. Mikosh and Starica (1999) use a GARCH model to detect structural changes in financial data. Diebold and Inou (1999) show analytically that long memory in processes (for instance the volatility process) may be due to the presence of structural breaks. Jackwerth and Rubinstein (1996) analyze option prices on the S&P 500 index and observe different probability distributions for implied binomial trees before and after the crash of 1987. Davies (2003) estimates piecewise constant volatility models. There the piecewise constant volatility function is selected because of its simplicity. Davies analyzes a similar data set as we do using an entirely different method for estimating the number and location of breakpoints. He finds 64 breakpoints over a 30 year period for the S&P 500 index. Bai and Perron (1998, 2001) investigate breakpoints in a multiple linear regressions model which is applied to interest rates over a 25 year period.

The effect of news on financial processes has been studied extensively. Berry and Howe (1994) show that the number of news items released by the Reuters News Service is significantly related to the volume but insignificantly related to the volatility of the S&P 500. See also the work of Cutler, Poterba and Summers (1989) where both political and economic news items are linked to price fluctuations. Contrary to the two above cited papers Ederington and Lee (1993) find that large fluctuations in financial markets may be coupled to significant economic news events. Andersen and Bollerslev (1998) use high frequency data to investigate the effect of economic news on foreign exchange rates, US\$ vs DM. Macro-economic news can clearly be linked to time-points of high volatility. However, they find little evidence of lasting effects. The main impact has disappeared after ten to twenty minutes.

We mention a possible application of our research. Duan, Popova and Ritchken (2002) derive option pricing formulae in a framework where volatility switches from regime to regime. In their setup the number of states is finite. See also the work of Buffington and Elliott (2002) for the application of option pricing under regime switching models. Our empirical analysis may provide some information on the parameters of such a model. Since we ob-

serve regimes which last for several months or even years, such a model may be useful for pricing long term options.

The present paper is organized as follows. Section 2 describes our data set. In section 3 we construct optimal piecewise linear approximations to the graph of \hat{q} with a prescribed number of breakpoints. In section 4 a simple probabilistic model is introduced. This allows us to estimate the number of breakpoints and to construct confidence intervals for the breakpoints. We find 26 breakpoints. In section 5 we try to link these breakpoints to news items indicating severe financial or economic shocks. The paper closes with a short discussion and a comparison with the breakpoints found in the integrated squared volatility of a small European stock index, the Dutch AEX.

2 Data

In this section we give a brief description of the data. Our data set consists of intraday future prices on the S&P 500, traded at the Chicago Mercantile Exchange over the period January 2, 1998 - August 31, 2001. We always use the future contract with the shortest time horizon. There are four expiration months: March, June, September, and December. We use future prices rather than the cash index in order to avoid the effects of non-synchronous trading, see Dacorogna et al. (2001). The bid-ask effect disappears for two minutes intervals. This determines the choice of the length of our interval. The selection from our data set consists of approximately 650,000 observations. In total we have 3430 trading days. The daily quadratic variation q_d is obtained by summing the squared log increments, corrected for a linear drift. (The drift term is so small that it could have been omitted without affecting our results.) The time-change \hat{q} , plotted in figure 1, is given by

$$\hat{q}(d) = \sum_{i=1}^d q_i, \quad d = 1, \dots, 3430.$$

In Peters and de Vilder (2002a) it is shown that the hypothesis that the time-changed S&P 500 is a Brownian motion can not be rejected. In the latter paper one may find more details on our data set, and the basic model.




Figure 1: The time-change \hat{q} . The points b mark visually detected breakpoints and b marks volatility bursts.

3 Piecewise linear approximation

This section treats a simple analytic problem. Suppose we have a discrete time function $\hat{q} : \{0, \dots, D\} \rightarrow \mathbb{R}$ and a positive integer $m \geq 1$. We are asked to determine the best piecewise linear approximation l_m to \hat{q} that agrees with \hat{q} in the m breakpoints, b_1, \dots, b_m and in the endpoints $b_0 = 0, b_{m+1} = D$. For elements i, j in $\{0, \dots, D\}$ with $i < j$ we shall write

$$[i, j] = \{i, i + 1, \dots, j - 1, j\} \quad (i, j) = \{i + 1, i + 2, \dots, j - 1\}$$

for the closed and open integer intervals with end points i and j . The breakpoints divide the interval $[0, D]$ into $m + 1$ subintervals $J_i = [b_i, b_{i+1}]$, for $i = 0, \dots, m$.

In our case $D = 3430$ (the number of trading days from January 1988 up to August 2001), and for $m = 20$ there are

$$\binom{D-1}{m} \approx 2 \cdot 10^{52}$$

piecewise linear approximations with m breakpoints. An optimization procedure over all possible partitions of the set $[0, D]$ by m breakpoints is not

feasible. Even if one imposes the condition that the distance between successive breakpoints equals at least 50 there still are far too many cases to consider.

It has been suggested to use a dynamic approach, the sample splitting method, see Chong (2001). This simple method yields a piecewise linear approximation in $O(mD)$ steps. In this method one first determines the best approximation with one breakpoint. Now proceed inductively. If the approximation l_k has been found, then for each of the intervals J_0, \dots, J_k , find the best approximation with one breakpoint. Finally, pick that interval which yields the largest decrease in the error, where the error is defined as the distance between the original function and its piecewise linear approximation.

Unfortunately, this method does not always give the best piecewise linear approximation. This flaw is demonstrated in the next example.

Example 1 Consider the piecewise linear function in figure 2. Distance being measured in the L^2 norm, the first optimal point and hence the first breakpoint is 50. The second division points are 40 and 60. So for $m = 2$ one obtains either the partition $\{40, 50\}$ or $\{50, 60\}$, but not $\{40, 60\}$.




Figure 2: An example of a piecewise linear function with two breakpoints.

We shall now present a method that takes $O(D^3)$ steps. This method is based on the Suboptimization Principle in Fisher (1958), page 795, which states: If l_m is the optimal approximation to \hat{q} on $[0, D]$ with m breakpoints b_1, \dots, b_m then the restriction of l_m to $[0, b_m]$ is the optimal approximation to \hat{q} on $[0, b_m]$ with $m - 1$ breakpoints.

In order to determine the best piecewise linear approximations we need a measure ε for the error of the approximation. There are many candidates, that strictly speaking need not be distance measures:

$$\int |q - l_m| dx, \quad \int (q - l_m)^2 dx, \quad \left(\int (q - l_m)^2 dx \right)^{\frac{1}{2}}, \quad \sup (q - l_m)^2.$$

Given the piecewise linear approximation l_m with breakpoints b_1, \dots, b_m let $\varepsilon_i = \varepsilon[b_i, b_{i+1}]$ for $i = 0, \dots, m$ denote the error of the (linear) approximation l_m over the interval $[b_i, b_{i+1}]$. We shall take an axiomatic approach, introducing three reasonable conditions the error function has to satisfy. In the next subsection we prove that it is possible to obtain the optimal partition for any measure ε which satisfies our three conditions.

1) We assume that the error ε over the whole interval $[0, D]$ is a function of the $m + 1$ error terms $\varepsilon_0, \dots, \varepsilon_m$ over the intervals $[b_i, b_{i+1}]$:

$$\varepsilon = \gamma(\varepsilon_0, \dots, \varepsilon_m). \quad (1)$$

2) In view of the suboptimization principle we shall assume:

$$\gamma(\varepsilon_0, \dots, \varepsilon_{k+1}) = \gamma(\gamma(\varepsilon_0, \dots, \varepsilon_k), \varepsilon_{k+1}), \quad k \geq 1. \quad (2)$$

So it suffices to define $\gamma(\varepsilon_0, \varepsilon_1)$ for two variables $\varepsilon_0, \varepsilon_1 \geq 0$.

3) We restrict attention to error measures for which $\gamma(\varepsilon_0, \varepsilon_1)$ is increasing in ε_0 :

$$\varepsilon_0 < \varepsilon'_0 \implies \gamma(\varepsilon_0, \varepsilon_1) \leq \gamma(\varepsilon'_0, \varepsilon_1). \quad (3)$$

In the four examples above $\gamma(\varepsilon_0, \varepsilon_1)$ has the form

$$\varepsilon_0 + \varepsilon_1, \quad \varepsilon_0 \vee \varepsilon_1, \quad \sqrt{\varepsilon_0^2 + \varepsilon_1^2}, \quad \varepsilon_0 \vee \varepsilon_1,$$

respectively.

3.1 The algorithm

In each of the four examples mentioned above one may use the following procedure to obtain the optimal approximation with m breakpoints for all $m \geq 1$.

1) First compute the symmetric matrix E with entries

$$e_{ij} = \varepsilon[i, j].$$

This takes $O(D^3)$ steps. Note that the diagonal entries are zero, and so are the entries e_{ii+1} .

2) Define the functions e_1 and z_1 on $[0, D]$ by

$$\begin{aligned} e_1(n) &= \min_{0 \leq l \leq n} \gamma(e_{0l}, e_{ln}), \\ z_1(n) &= \operatorname{argmin}_{0 \leq l \leq n} \gamma(e_{0l}, e_{ln}) \end{aligned}$$

where γ is the function in (1). Then $e_1(n)$ is the error term for the best continuous piecewise linear approximation to \hat{q} on the interval $[0, n]$ with one breakpoint, and the breakpoint is $z_1(n)$.

3) Similarly define the functions e_k and z_k for $k = 2, \dots, m$ by

$$e_k(n) = \min_{0 \leq l \leq n} \gamma(e_{k-1}(l), e_{ln}), \quad (4)$$

$$z_k(n) = \operatorname{argmin}_{0 \leq l \leq n} \gamma(e_{k-1}(l), e_{ln}). \quad (5)$$

Then $e_k(n)$ is the error term for the best continuous piecewise linear approximation to \hat{q} on the interval $[0, n]$ with k breakpoints. In total $O(mD^2)$ steps are needed to define the function ε_k and z_k for $k = 1, \dots, m$.

Proposition 1 *Suppose the error term satisfies (1), (2) and (3). Then $e_m(D)$ as defined in (4) is the minimal approximation error for a continuous piecewise linear approximation with m breakpoints. The breakpoints b_1, \dots, b_m are obtained by the backward recursion*

$$b_m = z_m(D), \quad b_k = z_k(b_{k+1}), \quad k = m-1, \dots, 1,$$

with z_k defined in (5).

Proof. The proof is by induction. For $m = 1$ all partitions with 1 breakpoint are considered. The procedure returns the optimal approximation: breakpoint $z_1(D)$ and error term $e_1(D)$. Now suppose the optimal approximations on $[0, d]$ have been obtained for $m = n$ breakpoints and $d = 1, \dots, D$. Since the error measure satisfies (2) and (3) the minimal approximation error for $m = n + 1$ breakpoints can be found by

$$\min_d \gamma(e_n(d), \varepsilon(d, D))$$

with $e_n(d)$ the approximation error of the optimal piecewise linear approximation on the interval $[0, d]$. Here we use Fisher's suboptimization principle. This defines $e_{n+1}(D)$ and determines the $n + 1$ breakpoints of the optimal linear approximation. In the same manner we find the optimal piecewise linear approximation with $n + 1$ breakpoints on $[0, d]$ for $d < D$, and its error. \square

Because of the stochastic background to our problem we shall use the error measure

$$\varepsilon = \sum_{d=0}^D (\hat{q}(d) - l_m(d))^2. \quad (6)$$




Figure 3: The approximation error ε on a logarithmic scale for different number of breakpoints $m = 1, \dots, 50$.

Figure 3 depicts the decrease of the error for $m = 1, \dots, 50$. Figure 4 exhibits the position of the breakpoints as the number of breakpoints increases from one to fifty. After six breakpoints there is an abrupt change in the configuration and after 21 breakpoints as well. The other changes are gradual. The positions of some breakpoints tend to be more stable than others. On the whole the stability of the breakpoints is striking.

From the figures 3 and 4 it is not clear which partition we should choose. In the next section we introduce a stochastic model. This will allow us to distinguish between spurious and true breakpoints.




Figure 4: The location of the breakpoints for $m = 1, \dots, 50$. At the bottom the location of the breakpoints after pruning as described in section 4.2.

4 A probabilistic model

In the previous section we presented an algorithm for finding the optimal piecewise linear approximation with m breakpoints. This section is concerned with the problem of determining the number of breakpoints and finding confidence intervals around their location. It is ultimately our aim to link breakpoints to news items indicating changes in the economic climate. We shall use the quadratic error measure (6) to quantify the difference between the approximation and the actual function. The number of breakpoints can not be estimated by minimizing the approximation error over all possible partitions of breakpoints. This would yield $m = 3430$ breakpoints.

In order to determine the number of breakpoints a probabilistic model is introduced. We shall assume that the sample function \hat{q} of the quadratic variation, which we are trying to estimate, is the sum of a continuous piecewise linear function f and a sample function of Brownian motion. Very little is known about the quadratic variation process. Our choice of Brownian motion to model the error term is motivated by the absence of good alternatives. We do not assume that the increments are normally distributed nor that they are independent. Under suitable mixing conditions many discrete time processes can be modelled asymptotically by a Brownian motion. Therefore the use of the Brownian motion model for the error process in our situation is defensible. In this model the function \hat{q} is an estimate of the sample function of the process Q

$$Q = f + \delta B \tag{7}$$

where B is a standard Brownian motion restricted to the integers and f is piecewise linear.

The Brownian motion δB contains the statistical error made in estimating the quadratic variation on the basis of the 650 000 high frequency data points. This statistical error is cumulative.

In addition to this statistical error there are random fluctuations in the volatility itself. Daily volatility σ_d is known to vary considerably from day to day. This variation yields random fluctuations in the daily increments σ_d^2 of the quadratic variation. The random daily volatilities are correlated but their overall effect on the quadratic variation, over extended periods of time, may still be modelled by a Brownian motion.

Let G be the continuous piecewise linear process which agrees with the process Q in the breakpoints of f in (7). Then the difference $Q - G$ is a con-

catenation of independent Brownian bridges. It is the sample function g of the process G which we are trying to estimate.

Define a *regime* as the interval between two successive breakpoints. The slope of \hat{q} over this interval will be called the *activity* α . This is the average of the squared volatility during the regime. The squared volatility fluctuates around this average value α . The sum of the deviations is cumulative and constitutes the second component of the Brownian motion δB .

In short, the process δB is the sum of two processes, one due to the fact that our data set is finite, the other an inherent part of the quadratic variation process. We emphasize that there is little hard evidence for our model. The model is selected because of its simplicity and general applicability.

We shall first describe the error process. The subsection thereafter focusses on the number of breakpoints by making distinction between spurious and non-spurious breakpoints. Section 4.3 determines confidence intervals for the breakpoints and for the activity.

4.1 The error process E_A

For any finite subset A of $[0, D]$ we shall define an error process E_A . Suppose $A = \{a_1, \dots, a_m\} \subset (i, j)$ with $i = a_0 < a_1 < \dots < a_m < a_{m+1} = j$. With A we associate the piecewise linear process L_A on $[i, j]$ which agrees with Q in the points a_0, \dots, a_{m+1} . The error process E_A is the difference between L_A and Q . If A contains the breakpoints in (i, j) then E_A is a concatenation of Brownian bridges. In order to have an explicit description of the error process E_A we introduce independent standard Brownian bridges B_0^0, \dots, B_m^0 on $[0, 1]$. (Recall that the standard Brownian bridge B^0 may be obtained from the standard Brownian motion B by setting $B^0(t) = B(t) - tB(1)$.) Set $\Delta_k = a_{k+1} - a_k$ for $k = 0, \dots, m$. For $d \in [i, j]$ write

$$Q(d) = L_A(d) + \delta \sqrt{\Delta_k} B_k^0(u), \quad d = a_k + u\Delta_k, \quad k = 0, \dots, m.$$

Here δ^2 is the variance of the noise process in (7). The process $E_A = Q - L_A$ will be used to investigate the breakpoints of the function \hat{q} .

4.2 The number of breakpoints

Our first task is to determine the number of breakpoints in the graph of \hat{q} in figure 1. In the previous section we presented an algorithm to find the

location of the breakpoints given the number of breakpoints. For the error measure in (6) an increase in the number of breakpoints will not always decrease the approximation error. But in general more breakpoints should give a better approximation, as it does for our dataset when m increases from 1 to 50. See figure 3.

For large m it may happen that some of the m breakpoints determined by our algorithm in section 3 do not correspond to breakpoints of the piecewise linear function f in our model (7) but are due to chance fluctuations in the noise term δB . Such breakpoints will be called *spurious*. Spurious breakpoints form the subject of the second subsection below. Subsequently we describe a pruning operation which will be used to determine the final set of breakpoints.

4.2.1 The parameter δ

We start by determining a first estimate of the parameter δ . In our case we do *not* assume that the increments of the error process are i.i.d. Gaussian variables. Figure 5 shows that the time series of the daily increments of the quadratic variation is not a realization of the sum of a piecewise constant function and Gaussian noise. In fact, autocorrelation in the error process




Figure 5: The daily increase of the quadratic variation.

occurs at lags 1 and 2. To filter out the autocorrelation we define

$$\delta_k^2 = \frac{1}{D} \sum_{d=1}^{\lfloor D/k \rfloor} (e_m((d-1)k) - e_m(dk))^2,$$

which may be interpreted as estimates of the quadratic variation of $\hat{q} - l_m$ divided by D based on increments of length k . Take $m = 30$. Due to autocorrelations δ_k is small for $k = 1$. For $k > 1$, δ_k fluctuates around the mean value $9.5 \cdot 10^{-4}$. Since we are interested in the variation on a large time scale we take this mean value as an initial estimate for δ . Figure 6 plots the quadratic variation of the error process calculated from 2 day increments. Observe that the figure does not support the hypothesis that δ is constant over time.




Figure 6: The quadratic variation of the error function based on intervals of two days. The vertical jumps correspond to the wide confidence intervals of the activity in figure 13.

4.2.2 Spurious breakpoints

A natural way to decide whether an estimated breakpoint is spurious is to compute the probability that it is spurious. This is done by estimating a breakpoint in a Brownian bridge and determining the decrease in the 'approximation error'. The probability that the breakpoint is spurious is then computed from the distribution of this decrease in the approximation error.

Suppose we are given a set of possible breakpoints $\{b_1, \dots, b_n\}$ that contains the set of true breakpoints. Consider the point b_k . The standardized reduction r_k of the approximation error due to the breakpoint b_k is given by

$$r_k = \frac{\int_{b_{k-1}}^{b_{k+1}} (q(t) - l_k(t))^2 dt - (\varepsilon_{k-1} + \varepsilon_k)}{\delta^2 (b_{k+1} - b_{k-1})^2} \quad (8)$$

where $l_k(t)$ is the linear function through the two points $(b_{k-1}, q(b_{k-1}))$ and $(b_{k+1}, q(b_{k+1}))$, for $k = 1, \dots, n$. If b_k is spurious then the error process on the interval $[b_{k-1}, b_{k+1}]$ is of the form

$$E(d) = \delta \sqrt{b_{k+1} - b_{k-1}} B^0 \left(\frac{d - b_{k-1}}{b_{k+1} - b_{k-1}} \right), \quad d \in [b_{k-1}, b_{k+1}]. \quad (9)$$

Define the random variable R by

$$R = \int_0^1 (B^0(t))^2 dt - \int_0^1 (B^0(t) - L_1(t))^2 dt,$$

where $L_1(t)$ is the piecewise linear approximation with one breakpoint T to the Brownian bridge $B^0(t)$ on $[0, 1]$.

If b_k is spurious, r_k may be viewed as a realization of R conditionally on $T = t_k$ where t_k denotes the standardized location of the breakpoint b_k ,

$$t_k = \frac{b_k - b_{k-1}}{b_{k+1} - b_{k-1}}.$$

Testing the null hypothesis that b_k is spurious we may reject for large values of R . Given r_k the corresponding p-value is

$$P_k = P((R|T = t_k) > r_k). \quad (10)$$

This is the probability, under the null hypothesis, to find a more extreme value than r_k for R given $T = t_k$.

We have not been able to derive an explicit expression for the bivariate distribution of (R, T) . Therefore this random vector has been simulated. The reductions r_i are compared to simulated realizations of the conditional variables $(R|T \in [t_i - \eta, t_i + \eta])$, with $\eta = 0.01$. Figure 7 shows the estimate of the marginal density of T based on one million simulations. Fortunately T is approximately uniformly distributed on the interval $[0, 1]$. This simplifies the task of finding realizations of the stochastic variable $(R|T \in [t_i - \eta, t_i + \eta])$.




Figure 7: Simulated density of the location of a spurious breakpoint T .

4.2.3 The pruning operation

The algorithm in section 3 returns the optimal approximation for a given number of breakpoints. It may happen that one obtains spurious breakpoints instead of real breakpoints. In particular this may happen if the piecewise linear function contains a long regime, as is shown in the example below.

Example 2 *The process depicted in figure 4.8 is the sum of a piecewise linear function (dotted) and a concatenation of three Brownian bridges. The true*




Figure 8: An example of a sum of a piecewise linear function, with two breakpoints located such that one long regime and two short regimes are present, and a concatenation of three Brownian bridges.

breakpoints are 80 and 90. If we estimate a piecewise linear approximation with 2 breakpoints we find the partition (42,71). The probabilities are $P_1 = 0.17$ and $P_2 = 0.07$. The value of $P_1 = 0.17$ indicates that there is a chance of 1 in 6 that the first breakpoint is spurious. If we estimate a piecewise linear approximation with 3 breakpoints we find (40,75,93) with corresponding probabilities (0.30,0.00,0.063). These numbers suggest that 75 and 93 are real breakpoints and 40 is spurious.

The example shows that a long regime may affect the finding of breakpoints for regimes of smaller length. Therefore we shall perform a pruning operation which removes spurious breakpoints from a given set of possible breakpoints. Our strategy will be as follows. First the number of breakpoints N is set at a large level, larger than the anticipated number m of real breakpoints. Given this N , the breakpoints b_1, \dots, b_N are determined and subsequently we try to remove the spurious breakpoints by a cleaning operation. This pruning operation is akin to the pruning procedure in Breiman et al. (1984). The precise algorithm we use is described below.

Step 1. Choose a level β , say $\beta = 0.05$.

Step 2. Estimate the piecewise linear approximation l_N . This gives us a set A_N of N breakpoints. At this point we assume that all the non-spurious breakpoints are in the set A_N . With each breakpoint b_i , $i = 1, \dots, N$, we associate the probability P_i that b_i is spurious. This probability is given by equation (10).

The pruning is done recursively by deleting the breakpoint with the highest P -value and recalculating the two nearest breakpoints. Suppose b_k is deleted. Then b_{k-1} and b_{k+1} are recalculated. Since P_{k-2} and P_{k-1} depend on b_{k-1} , and P_{k+1} and P_{k+2} depend on b_{k+1} we may have to recalculate in total four probabilities. Formally, given the pruning level β and the N breakpoints with associated probabilities P_1, \dots, P_N we prune as follows. If $\max(P_1, \dots, P_N) \leq \beta$ then we are done. Otherwise, we perform three steps:

Step 3. Set $k = \operatorname{argmax}_i \{P_i\}$. If P_k exceeds the level β then the breakpoint b_k is regarded as spurious and is removed from the set A_N of breakpoints.

Step 4. The breakpoint(s) b_{k-1}, b_{k+1} (b_2 if $k = 1$, b_{N-1} if $k = N$) are re-estimated by approximating the function on the interval $[b_{k-2}, b_{k+2}]$ ($[0, b_2]$, $[b_{N-2}, D]$) by a piecewise linear approximation with 2 (1) breakpoints. (The correction is performed on those two (one) breakpoints since their location is most strongly influenced by the removal of b_k .)

Step 5. The four (or three or two) probabilities P_j associated with the two (or one) points on either side of b_k are re-calculated. Each cycle decreases the number of breakpoints by one. The cycle is repeated until all P_i lie below the level β .

Applying this procedure to the function in example 2 results in the breakpoints (81,93) with corresponding probabilities (0.033,0.046). For the graph of \hat{q} we choose $N = 50$ and $\beta = 0.05$. We end up with a set A_{26} consisting of 26 breakpoints. The procedure seems to be quite stable. The same set of breakpoints is obtained on starting with any value of N between 45 and 52. Figure 9 depicts the partition of breakpoints that would have been obtained when starting with $N = 35, 40, \dots, 100$. For $N \geq 45$ the configuration hardly changes. The additional breakpoints that are obtained between $50 < N < 100$ lie very close (one day) to one of the elements of the set of breakpoints that corresponds to $N = 50$. A pair of breakpoints which lie close to each other is replaced by a triplet of breakpoints.




Figure 9: The location of the breakpoints after pruning. Left the number of breakpoints at te beginning. Right the number of breakpoints after pruning.

Since our arguments for a Brownian error process are based on longer periods than a day we stick to the set of 26 breakpoints which is obtained when starting with $N = 50$. The location of these 26 breakpoints is also given at the bottom of figure 4.

Note that the approximation error given by $l_{A_{26}}$ ($7.48 \cdot 10^{-4}$) is larger than the minimal approximation error for 26 breakpoints ($4.30 \cdot 10^{-4}$). (A spurious breakpoint in a long regime may decrease the error term more than a real

breakpoint in a small interval.) The piecewise linear approximation to \hat{q} based on these 26 breakpoints is given in figure 10.




Figure 10: The piecewise linear approximation to the time-change \hat{q} with 26 breakpoints.

The estimate of the parameter δ , $9.5 \cdot 10^{-5}$, was obtained from the error function e_{30} . If we estimate the parameter δ from the error function that follows from the function $l_{A_{26}}$, we find $9.7 \cdot 10^{-5}$. We would have obtained the same breakpoints if we had used $\delta = 9.7 \cdot 10^{-5}$.

We have also investigated an error function which depends on the length of the regimes, replacing the terms $\varepsilon_i = \sum_{d=b_i}^{b_{i+1}} (q(d) - l_m(d))^2$ in $\varepsilon = \varepsilon_0 + \dots + \varepsilon_m$ by $\varepsilon_i / (b_{i+1} - b_i)$, see equation 6. This alternative error function has an unexpected side effect: long intervals are not sub-divided.

4.3 Confidence intervals

In this section we describe two procedures for defining confidence intervals for the breakpoints b_1, \dots, b_{26} which were found in the previous subsection. We also give confidence intervals for the activity level.

4.3.1 Maxima on either side of the breakpoint

The local behaviour at a breakpoint may be modelled by the process

$$X(t) = \begin{cases} at + \delta B_1(t) & t \geq 0 \\ bt + \delta B_2(-t) & t < 0 \end{cases} \quad (11)$$

with B_1 and B_2 independent Brownian motions and $a \neq b$. Define the process $Z(t)$

$$Z(t) = B(t) - |t| \quad t \in \mathbb{R} \quad (12)$$

where $B(t)$ is a two-sided standard Brownian motion passing through the origin. The processes on the positive and negative time axes are independent. Note the scaling property:

$$cZ(t/a) = cB(t/a) - |ct/a| \stackrel{d}{=} (c/\sqrt{a})B(t) - |ct/a|.$$

If we subtract the average slope $(a+b)/2t$ from $X(t)$ then we may model the local behaviour around a breakpoint in equation (11) by the process $Z(t)$ by a suitable scaling procedure.

Instead of the process Z we shall use a process Y on the interval $[-100, 100]$ of the form $Y(t) = -|t| + B^*(t)$ where $B^*(t)$ consists of two properly scaled independent Brownian bridges on the intervals $[-100, 0]$ and $[0, 100]$. For each realization $Y(\omega)$ we compute the optimal piecewise linear approximation $L_1(\omega)$ on $[-100, 100]$ with one breakpoint $T(\omega)$. Let T_1 denote the position of the maximum of Y over the interval $[-100, 100]$. This variable may lie on either side of T . Let T_2 be the position of the maximum of Y on the other side of T . So T lies between T_1 and T_2 . In many cases the real breakpoint, the origin, will also lie between T_1 and T_2 . We shall compute the probability P_0 that the origin lies between T_1 and T_2 . Ten thousand simulations give the value

$$P_0 = 0.49. \quad (13)$$

So in half of the cases the confidence interval with end points T_1 and T_2 will contain the true breakpoint. Since the confidence interval $[T_1, T_2]$ is defined geometrically the points T_1^i and T_2^i can be read off from the graph of \hat{q} around the point b_i by constructing a line through $(b_i, \hat{q}(b_i))$ whose slope is the average of the slopes of l_m on either side of the breakpoint and determining the maxima of \hat{q} above this line. Figure 4.3.1 illustrates how the points T_1 and T_2 are obtained.

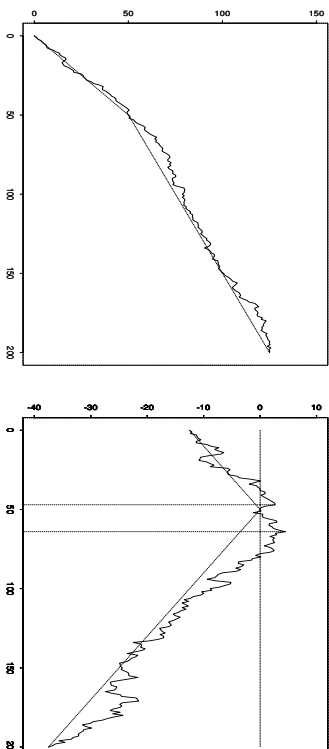


Figure 11: Locating the points T_1 and T_2 as discussed in section 4.3. From the graph in picture on the left we subtract the line with a slope equal to the average of the two slopes to the left and the right of the estimated breakpoint, $T=50$. This yields the picture on the right. It is easy to determine the points $T_1 = 64$ and $T_2 = 47$.

4.3.2 Likelihood functions

Let $i \in [1, 26]$. For $\theta \in \Theta = (b_{i-1}, b_{i+1})$ let l_θ be the piecewise linear function which agrees with \hat{q} in the points b_{i-1}, θ, b_{i+1} , and let $X = E_\theta$ be the associated error process on $[b_{i-1}, b_{i+1}]$ as described in section 4.1. For each realization of X let b_X be the breakpoint which minimizes the error. Set $p(\theta) = p_\theta(b_i) = P(b_X = b_i)$. Choose $\beta \in (0, 1)$, say $\beta = 1/10$, and define the confidence set

$$S_i = \{\theta \in \Theta \mid p_\theta(b_i) > \beta \sup_{\eta} p_\eta(b_i)\}.$$

The set S_i need not be an interval.

The probabilities $p_\theta(b_i)$ have been computed by simulating 10,000 Brownian bridges.

In order to give an indication of the outcome of this procedure we have plotted the likelihood function $\theta \mapsto p_\theta(b_{18})$. See figure 12.

Table 1 gives two confidence intervals for each breakpoint. One determined by the likelihood, the other based on the positions T_1 and T_2 of the maxima described in the first part of this section.




Figure 12: Simulated likelihood function of the 18th breakpoint.

4.3.3 Confidence intervals for the activity

It is not very hard to create confidence intervals for the activity in our model. The estimator of the activity in regime i is given by:

$$\alpha_i = 10^6 \frac{Q(b_i) - Q(b_{i-1})}{b_i - b_{i-1}} \quad (14)$$

The activity is blown up by the factor 10^6 for convenience. In that case the average activity over the entire period is 83.

Since the increment $\delta B(b_i) - \delta B(b_{i-1})$ of a Brownian motion is normally distributed with variance $\delta^2(b_i - b_{i-1})$, the variance of α_i is given by $10^{12} \delta^2 / (b_i - b_{i-1})$. However, as was observed in section 4.2 the variance δ^2 may not be constant over the whole period. See figure 6. Therefore, the factor δ^2 is replaced by the variance of the error term within a regime, δ_i^2 , which is calculated from two days intervals of the error function.

Figure 13 depicts the activity together with the borders of the confidence intervals. The 95% confidence intervals are obtained by adding/subtracting 1.96 standard deviations of α_i from the estimate of the activity in each regime.

5 Economic shocks

This section is concerned with linking breakpoints to news items.




Figure 13: The activity plotted on a logarithmic scale.

We set out trying to find a good piecewise linear approximation to the graph of \hat{q} in figure 1. Using Fisher's suboptimization principle it is possible with the current computers to obtain the optimal piecewise linear approximation with m breakpoints, for $m = 1, \dots, 100$. The L^2 norm was used to measure the error. In order to find the final number of breakpoints we have used a pruning technique. The resulting breakpoints form a very stable configuration. See figure 9.

The pruning operation deletes all breakpoints which have a probability of $\beta = 0.05$ or more of being spurious. So in our set of 26 breakpoints we may expect one or two to be false. For each of the 26 breakpoints we have given a confidence interval with confidence level 0.95. So we may expect one or two of the breakpoints to lie outside its confidence interval.

These expected numbers hold if the underlying model in (7) is valid. This model presumes the noise term to be a Brownian motion. For short intervals this assumption is definitely not correct because of autocorrelations, and because we apply a continuous model to discrete data. The difficulty is well illustrated by the graph of the daily increases in financial time in figure 5 which shows clusters of very high peaks rather than a piecewise constant function perturbed by Gaussian noise. That also is the reason why we have opted

for 26 breakpoints above, rather than 27, 28, or 29; see figure 9. Moreover, the assumption that the scale parameter δ of the noise process is constant is disputable. In figure 6 we present the graph of the quadratic variation of this noise process. A sizeable proportion of the increase is due to a few very short time intervals.

In spite of the rough methods used in analyzing the graph of the quadratic variation in figure 1 a visual inspection shows that the piecewise linear function in figure 10 yields a good fit. Figure 13 shows the 27 activity levels in logarithmic scale with their confidence intervals. Recall that the activity is the slope of the quadratic variation \hat{q} and hence the average of the squared volatility over a regime.

It is our aim to connect news items to the breakpoints. These news items should represent economic or political shocks which are so severe that they have a long term effect on the volatility. We list ten such severe shocks during the period under observation, January 1988 until August 2001. This list by nature is subjective. The reader is invited to change it at choice. The period of almost 14 years include four US-presidential elections, the first Gulf War, the breakdown of the communist system, and the introduction of the Euro. It is also the period of the ICT explosion. Here is our list of ten shocks:

- 1) October 13, 1989: October 1989 mini crash.
- 2) November 9, 1989: The Fall of the Berlin Wall.
- 3) August 2, 1990: Iraq invades Kuwait.
- 4) March 3, 1991: The First Gulf War ends.
- 5) December 6, 1996: Greenspan: 'Irrational exuberance'
- 6) July 2, 1997: Devaluation of the Thai Bath. Start of the Asia crisis.
- 7) July 28, 1998: Lewinsky agrees to testify. Impeachment becomes possible.
- 8) August 17, 1998: Devaluation of the Ruble. Russian default.
- 9) September 23, 1998: Collapse and rescue of the Long Term Capital Management (LTCM) Hedge Fund.
- 10) January 13, 1999: Devaluation Brazilian Real: Samba Crisis.

The fundamental question of what one means by news will not be treated here. We have two main problems concerning the nature of news items. First, one has to distinguish between positive and negative news and secondly, one has to establish the importance or intensity of the news.

In economics in general it is difficult to define what one means by positive or negative news. Events have no absolute meaning. The general state

of the market plays a role in the interpretation. For example, a decrease in the inflation rate may be interpreted both as positive and as negative news. If inflation is decreasing from a high level the news will be interpreted as positive. However, if there is a danger of deflation the news may be interpreted as negative. The 'same' news may have a different meaning in a different context.

Expectations of market participants also play a role in the way the market treats news events. News which is in line with leading expectations will not have a large impact on market movements.

We have looked through the pages of the New York Times (NYT) and the Wall Street Journal (WSJ) over the periods around the breakpoints suggested by the confidence intervals and the two maxima, bearing in mind the points made above. Sometimes the news is clear, but often it is hard to determine the intensity of the impact and to evaluate the effect of a news item on volatility. This holds in particular for political news. News is ambiguous. It is difficult to decide to what extent news items confirm expectations. News items on the same page of a newspaper may point in different directions.

5.1 Headlines of the Wall Street Journal and the New York Times

Below we give a short description of the major news events that we found in the NYT and WSJ.

The precise time of a news event may be important. A news shock on Monday May 11, 8 a.m. will be reported in the morning paper of Tuesday May 12 and will be reflected in a breakpoint on Friday May 9. The change in activity is visible only after May 9.

Table 2 contains the headlines of WSJ and NYT. Breakpoints which are clearly related to important news items are marked by an asterisk * both in the table and below. The numbers between the news items indicate the activity level during the corresponding regime. The initial regime has activity level 228. The activity over a regime is the average variance of the log returns over this period. This quantity has been multiplied by a million, see equation (14), to obtain more palatable numbers, which vary between twenty and five-hundred.

1. Strong indications that the U.S. economy is entering a period of weak growth. The Napm confidence index (01/31/1998) and leading economic indicators (02/02/1998) have declined. Markets expect Fed to lower interest rates (01/31/1998). This belief gets more widespread in the week prior to this regime change.

92

2. A stronger Dollar as the exchange rate increases on August 17, 18, and 24, 1988. The CPI increases by 0.4% after two increases by 0.3%. Inflation is accelerating, but not as fast as feared (0.5-0.6%).

45

The breakpoint 3 and 4 form a pair around the 1989 mini crash.

3*. Thursday October 5, 1989. The British interest rate moves up by one point to 15% and the German Bundesbank increases its key money rates by one point. On the domestic market on Friday October 13 the producer prices of September leapt up 0.9% (where 0.7% had been feared). The S&P drops by 9% but rebounds by 5% on Monday.

343

4. Thursday October 19, 1989 the consumer price increase over September is only 0.2%. Fears of a repeat of the October 1987 crash appear to be unfounded.

55

The breakpoints 5 and 6 also form a pair.

5*. On January 15 1990 the Wall Street Journal writes that producer prices have increased by 0.7%. They also write 'In the final quarter of 1989, economic growth is believed to have nearly come to a halt as consumer spending sagged, factories laid off workers and the nation's industrial production stagnated.'

165

26

6*. On Tuesday January 30, 1990 Greenspan describes the current economic slowdown as ‘temporary hesitation’. On the international scene Gorbachev does not rule out German reunion (January 30) and De Klerk promises to set Mandela free soon (February 2).

59

The breakpoints 7, 8, and 9 occur around the first Gulf War.

7*. On July 17, 1990 Iraq threatens to use force against Kuwait if it does not curb excess oil production. (The Iraqi invasion takes place two weeks later on August 2.) On the 18th the NYT reports that U.S. imports record 49.9% on oil. Domestic economic news: Homebuilders see recession and blame the saving crisis (07/19/1990 NYT). Inflation is higher than expected. On July 19 Greenspan announces that he does not cut the interest rate.

175

8. On October 18, 1990: Consumer prices rise 0.8% in September for the second month. Sharp increase in oil prices accounts for more than half the increase.

During the regime between breakpoint 8 and 9 there are two days of high volatility. On Wednesday January 9, 2001 the quadratic variation increases by 0.0016 (Baker-Aziz talks fail) and on Thursday January 17 by 0.0020 (start of the Gulf war). The factor 10^6 in equation (14) gives increases of 1600 and 2000 respectively for these days.

124

9*. The Gulf War ends on Sunday March 3, 1991. On March 12 the OPEC agrees on output cuts to raise prices. On March 13 Greenspan hints that the reduced threat of inflation might make it possible to give the economy an additional boost by lowering the interest rates.

53

10. This is the beginning of a long regime of low activity lasting four years. There are no clear news items, in the large confidence interval around breakpoint 10, that can be linked to this breakpoint, dd. April 28, 1992. Two periods start around this breakpoint. 1) Upcoming of the Democratic leader

27

Bill Clinton as challenger of Bush Sr. On April 8, 1992 Clinton wins New York. On April 29 Clinton and Bush win Pennsylvania. 2) This breakpoint may be marked as the beginning of the expansion of the U.S. economy. On April 17 Housing and Output figures give new sign to U.S. recovery. On May 6 the federal reserve reports that economic activity increases. On May 13 it is announced that the retail sales rose 0.9% in April.

31

11. Here ends the long period of low activity. The economy grows slowly but steadily. Unexpected rate cut in Europe on January 24, 1996. On January 18 U.S. blue chips show better earnings than expected. Activity increases on earning reports.

56

The breakpoints 12 and 13 form a pair. Over the intervening two week period the activity is high.

12*. The increased volatility seems to be triggered by a profit warning by Hewlett-Packard on Thursday July 11, 1996. Quote from WSJ July 12-13, 1996 (L. Bauman): 'The recent weakness of the group [technology] was accelerated by Wednesday's late-day announcement by computer maker Hewlett-Packard that it has suffered a significant, widespread decline in the growth of new orders during the current quarter. The outlook came on the heels of Motorola late Tuesday posting unexpectedly weak second-quarter earnings, which followed a string of disappointing earnings previews by a wide variety of technology companies'. On Monday July 15 the Nasdaq drops 3.9%.

195

13*. On Tuesday July 23, 1996 Greenspan signals that the Fed will not increase the interest rate. His words have a stabilizing influence on the market. He notes: 'The recent volatility in the stock market is the norm rather than the exception', see WSJ July 24, 1996. On July 24 Compaq publishes good second quarter results. So does IBM the next day. As a result the market calms down and the activity drops back to its level before July 11.

37

28

14*. During a dinner speech December 6, 1996, Greenspan tries to curb U.S. stock market's euphoria uttering the words 'Irrational exuberance'. The S&P 500 drops more than 2.7% in the first 30 minutes but it regains its previous level the next trading day. The activity increases.

90

15. There is turmoil in south-east Asia leading to lower exchange rates against the Dollar and higher interest rates. On August 13, 1997 Indonesia defends its Rupiah by intervention and increasing interest rates. On August 14 Indonesia floats the Rupiah. It drops 6%. On August 28 Asian stocks and currencies fall sharply. On October 8 Greenspan's comments on the U.S. economy chill the markets. He also raises the prospect of a rate rise. Activity increases.

163

16. Greenspan's words about the Asia crisis to the House banking committee on November 12, 1997 calms the market, as does the (expected) decision of the Fed not to raise the interest rates.

69

During the period August-October 1998 the activity is high. This may be due to the political insecurity caused by the Clinton-Lewinsky affair. The senate's acquittal of Clinton on the impeachment trial is on February 12, 1999.

17*. On Wednesday July 22, 1998 Greenspan warns that the Asia Crisis 'has shown no evidence of stabilization', and 'History tells us that there will be a correction of some significant dimension'. On July 28 Monica Lewinsky agrees to testify in turn for full immunity. Activity increases.

184

18*. On August 27, 1998 the S&P drops by 3.84%. The financial markets are influenced by sharp declines on the Russian markets and fears that Yeltsin may resign. (On August 25 the ruble drops by 9.2%.) Activity increases further.

408

29

19*. On Thursday October 15, 1998 the Fed cuts the federal fund rate from 5.25 to 5% citing unsettled markets. This is the first unscheduled cut since 1994. Stock markets rally world wide.

118

20. On March 6, 2000 stocks fall after Greenspan's speech. He still regards stocks as overpriced and warns that interest may rise. On March 21 the Fed indeed raises the interest rates by a quarter in order to dampen the economy.

262

21*. Macro economic figures indicate that the growth of the U.S. economy is slowing down. On Wednesday May 31 2000 the Napm index drops. On June 1 the U.S. manufacturing economy slowed considerably in May. On June 2 'the rise in joblessness delight U.S. markets'.

79

22*. Major Indices fall after profit warning of Kodak on September 26, 2000. Apple and UAL report on September 29 that earnings will not meet expectations. On October 3 the Fed leaves interest rates, as expected, unchanged despite concerns on inflation. Activity increases.

199

The breakpoints 23 and 24 form a pair. Over the intervening week the activity is very high. On the one day, January 3, 2001, the quadratic variation increases by 0.0018, approximately 0.7% of the total quadratic variation over 14 years. Visual inspection of the graph suggest that there is only one exceptional day between the 23rd regime (activity: 199) and the 25th regime (activity 102). Our procedure finds two breakpoints which are six days apart (rather than one day). This may be due to the particular form of the L^2 metric which we are using

23. No news between December 27, 2001 and January 2, 2001 that seems to be important. Yet the activity increases.

489

30

24*. On January 3 2001 the Fed surprises by cutting the interest rate from 5.5 to 5 percent. This is the first cut after six increases during the previous two years. The headline of WSJ the next day: 'Surprise Interest-Rate Cut by the Fed Catches Traders Off Guard, Helps to Lower Volatility'.

102

25*. On February 21, 2001, the Nasdaq falls 4.4% as analysts and companies cut their profit estimates. On February 22, 2001, the CPI increases by +0.6 percent. It is the largest increase in 10 months and twice as high as expected. The activity increases.

240

26*. On April 18th 2001 the Fed unexpectedly cuts the federal fund rate from 5 to 4.5 percent. Greenspan remarks that the economy is 'unacceptably weak'. The market interprets the rate cut as pro-active. This has a stabilizing effect on the stock market. The action of the Fed makes clear that the Fed is in control. This ends a high activity regime of average length.

96

Table 1: Confidence intervals of \hat{b}_i

The third column contains the confidence intervals of the estimates \hat{b}_i based on the likelihood (lh). The fourth column contains the date T_1^i and the fifth column T_2^i . If T_1^i coincide with b_i then T_2^i is not calculated.

i	Date b_i	LH	T_1	T_2
1	Feb. 8, 1988	Feb. 3 – Feb. 11	Jan. 21	Feb. 9
2	Aug. 25, 1988	Aug. 12 – Sep. 9	Aug. 16	Aug. 26
3	Oct. 9, 1989	Oct. 3 – Oct. 12	Oct. 12	Oct. 6
4	Oct. 24, 1989	Oct. 23 – Oct. 25	Oct. 17	Oct. 25
5	Jan. 11, 1990	Jan. 8 – Jan. 16	Jan. 11	-
6	Feb. 2, 1990	Jan. 29 – Feb. 8	Jan. 30	Feb. 5
7	Jul. 18, 1990	Jul. 12 – Jul. 23	Jul. 20	Jul. 17
8	Oct. 26, 1990	Oct. 24 – Oct. 31	Jan. 17	Oct. 18
9	Mar. 19, 1991	Mar. 11 – Mar. 28	Jan. 18	Mar. 20
10	Apr. 28, 1992	Mar. 26 – Jun. 4	Feb. 26	Apr. 29
11	Jan. 15, 1996	Dec. 5 – Feb. 12	Jan. 3	Feb. 8
12	Jul. 15, 1996	Jul. 9 – Jul. 19	Jul. 10	Jul. 19
13	Jul. 24, 1996	Jul. 18 – Jul. 30	Jul. 24	-
14	Dec. 2, 1996	Nov. 15 – Dec. 13	Dec. 2	-
15	Aug. 7, 1997	Aug. 1 – Aug. 12	Oct. 22	Aug. 6
16	Nov. 14, 1997	Nov. 10 – Nov. 20	Nov. 14	-
17	Jul. 28, 1998	Jul. 22 – Aug. 4	Jul. 22	Jul. 30
18	Aug. 25, 1998	Aug. 24 – Aug. 26	Aug. 26	Aug. 24
19	Oct. 22, 1998	Oct. 20 – Oct. 27	Oct. 15	Oct. 23
20	Mar. 10, 2000	Mar. 6 – Mar. 15	Mar. 29	Mar. 6
21	May 31, 2000	May 26 – Jun. 2	May 26	Jun. 1
22	Sep. 29, 2000	Sep. 26 – Oct. 3	Oct. 3	Sep. 28
23	Dec. 28, 2000	Dec. 27 – Dec. 29	Jan. 2	Dec. 27
24	Jan. 8, 2001	Jan. 8 – Jan. 9	Jan. 3	Jan. 9
25	Feb. 21, 2001	Feb. 20 – Feb. 22	Feb. 20	Feb. 22
26	Apr. 20, 2001	Apr. 18 – Apr. 24	Apr. 18	Apr. 23

Table 2: Headlines

Headlines of the WSJ and the NYT and the change in activity $\Delta\alpha$. The asterisk indicates that we believe that the regime switch may be explained by news.

No.	$\Delta\alpha$	Events
1	-136	02/01/1988, WSJ: Bonds analysts are bullish on sign of lower U.S. rates. 02/03/1988, WSJ: U.S. prime rate, bond yields hit lowest levels in months.
2	-47	08/18/1988, WSJ: U.S. dollar jumps to DM1.91 despite central bank action. 08/24/1988, WSJ: prices of U.S. bonds surge, rates ease on inflation data.
3*	+298	10/06/1989, WSJ: Eight European central banks raise rates after bundesbank boost discount. 10/16/1989, WSJ: U.S. Producer prices leapt 0.9% as surging energy costs rekindled inflation.
4	-287	10/20/1989, WSJ: U.S. Consumer prices rise by surprisingly small 0.2%.
5*	+110	01/13/1990, NYT: Producer prices up 4.8 percent in '89, most in 8 years.
6*	-106	01/31/1990, WSJ: Chance of U.S. recession has abated, Fed chief says. NYT: Wary gorbachev sees momentum toward Germany's reunification.
7*	+116	07/18/1990, NYT Iraq threatens Emirates and Kuwait in oil glut. 07/19/1990, WSJ: Greenspan signals rates to hold. NYT: Home builders see recession and blame the saving crisis.
8	-51	10/19/1990, NYT: Consumer prices rise 0.9 percent again.
9*	-71	03/04/1991, NYT: U.S. says Iraqi generals agree to demands 'On all matters'. 03/14/1991, NYT: Greenspan hints at cut in interest rates.
10	-22	04/08/1992 NYT: Survey gives Clinton the edge in New York democratic vote. 04/29/1992, NYT: Clinton and Bush capture primaries in Pennsylvania and look to a fall battle. 05/07/1996, NYT: Fed finds more signs of upturn. 05/14/1996, NYT: Retail sales rose 0.9 percent in April.
11	+25	01/16/1996, WSJ: Stocks fall sharply as worries continue about earnings. Technology stocks tumble ahead of Intel report. 01/18/1996, NYT: Fed survey sees 'modest pace of growth. 01/25/1996, WSJ: Industrials soar to record 5242.84 on earnings and merger. 01/26/1996, NYT: Disappointing earnings push stocks lower.
12*	+139	07/12/1996, WSJ: H-P earnings lead tech issues to losses on 7th-most-active day.
13*	-158	07/24/1996, NYT: Bond prices rally on Greenspan's comments and unexpected drop in retail sales figures. 07/25/1996, WSJ: Compaq numbers bolster tech stocks. 07/26/1996, NYT: I.B.M. soars as it beats quarterly expectations.
14*	+54	12/04/1996, WSJ: Fed chairman pops the big question: Is market too high? NYT: Greenspan asks a question and global markets wobble.
15	+73	08/29/1997, WSJ: Southeast Asian stocks, currencies fall sharply. 10/09/1997, NYT: Greenspan Cautions on inflation; change in tone a jolt to investors.
16	-94	11/14/1997, NYT: Crisis in Asia no big threat, Fed chief says. NYT: Iraq carries out threat to expel U.S. inspectors.
17*	+115	07/23/1998, NYT: Greenspan act II leaves little to cheer about on Wall St. 07/29/1998, WSJ: Lewinsky, granted immunity, will testify.
18*	+225	08/25/1998, NYT: German banks lead world in nail-biting over Russia. 08/27/1998, WSJ: Russia quits fight back to the ruble. NYT: Economic turmoil in Russia takes toll in Latin America.
19*	-290	10/16/1998 WSJ: Fed chairman cuts rates quarter point in a surprise move to shore up markets.
20	+144	WSJ: Greenspan views stocks as a threat. NYT: Reminder from Greenspan: expect interest rate to rise.
21*	-183	06/02/2000, NYT: Data indicates further slowing of economy. 06/05/2000, WSJ: Job losses suggest economy may be slowing.
22*	+120	09/27/2000. NYT: Major share indices fall as Kodak issues warning. 09/30/2000, NYT: Apple and UAL lead market downward as quarter ends.
23	+290	12/28/2002, WSJ: Dollar, Euro touch highs against Yen on more dismal numbers from Japan. 12/28/2000, NYT: Share prices move higher on hopes of easing by the Fed.
24*	-387	01/04/2001 WSJ: Surprise interest-rate cut by the Fed catches traders off guard, helps to lower volatility
25*	+138	02/22/2001, NYT: Inflation index jumps, thanks to energy costs
26*	-144	04/19/2001, WSJ: World markets cheer surprise rate reduction by U.S. Fed, hoping it marks turning point.

6 Discussion

Of our list of the ten major news shocks in section 5 only three fall inside the confidence interval of the breakpoint: October 1989 mini-crash, Greenspan's words 'Irrational exuberance' and the Clinton-Lewinsky affair. This is a disappointing score. The confidence intervals cover almost 10 percent of the period under consideration. So on average we should find one link.

Sixteen breakpoints have been coupled, in a more or less convincing way, to news events in the list above. These have been marked by an asterisk. The news events are unexpected rate cuts, remarks by Greenspan, inflation figures, profit warnings, and occasionally political events. These news events can not be regarded as the great economic shocks during this period. We conclude that the jumps in the activity level at the breakpoints are not caused by intense economic shocks. All we can say is that sometimes news items seem to precipitate such jumps.

Statements made in this section are by no means conclusive. We warn the reader that we do not understand what we see. It is our hope that this paper will inspire theorists to develop an economic theory which explains the phenomenon of a piecewise linear approximation.

6.1 Bird's eye view

We now take a bird's eye view of the behaviour of the activity of the S&P 500 futures stock index. Referring to figure 13 we find that the most striking feature is the long minimum in the center of the graph of the activity extending over more than one quarter of the total period. The graph also shows a certain degree of self similarity with the central part reflected to the left and right on a different level. The middle period corresponds to a sustainable growing economy. In the period on the left the economy is functioning below capacity. In the period on the right the economy is functioning above capacity and there is a danger of overheating.

Another striking feature of the figure are the four sharp peaks, each lasting only a few days. Each of the four peaks can unambiguously be linked to an unexpected high intensity news event. The first peak (b_3 and b_4) is initiated by unexpected high producer prices (inflation rate) and denotes the mini-crash of 1989. The second peak (b_5 and b_6) is also related to an unexpected increase in the producer prices. The third peak (b_{12} and b_{13}) coincides with very disappointing quarterly results of a number of leading technology stocks.

The fourth peak (b_{23} and b_{24}) is induced by an unexpected, and substantial, federal fund rate cut.

So much for the peaks. On the whole, regimes with high activity tend to be shorter with the exception of the central plateau around day 3000.




Figure 14: The activity in financial time plotted on a logarithmic scale.

Figure 14 records the activity levels in financial time. The change in the time scale yields a graph which at first sight bears no relation to graph in figure 13. If one leaves out the four peaks, then most of the remaining regimes have a length of approximately 0.010. The two exceptionally long regimes have lengths 0.029 and 0.041, which may be regarded as multiples of the basic length.

6.2 AEX

Before discussing the underlying model for the observed changes in activity we make a small detour to the Dutch stock exchange, the AEX. In Peters and de Vilder (2002b) a high frequency data set at 15 second intervals for the AEX index over the period May 1996–September 2000 is analyzed. Here too, the quadratic variation exhibits a piecewise linear structure. Starting with $N=50$ breakpoints we end up with a piecewise linear function with 19 breakpoints after pruning. For the S&P 500 we find 10 breakpoints in this period. For the sake of completeness we give in figure 15 the evolution of the location of the breakpoints as the number of breakpoints increases and the location of the breakpoints after pruning.




Figure 15: Above the location of the breakpoints of the AEX as the number of breakpoints increases and below the location after pruning.

The graph of the activity of the AEX together with the activity of the S&P 500 is plotted in figure 16.




Figure 16: The activity of the AEX plotted on a logarithmic scale. The dotted line is the activity of the S&P 500.

Note that the scales of the activity of the S&P 500 and the AEX in figure 16 are equal. There are four peaks for the AEX while there is only one for the S&P 500. Five of the breakpoints of the S&P 500, b_{14} , b_{15} , b_{17} , b_{18} , b_{19} , coincide with breakpoints of the AEX. At each of these five breakpoints the activity jumps in the same direction.

The confidence intervals of the ten breakpoints of the S&P 500 in the period May 1996–September 2000 cover 7.4 percent of the period. So one expects only one or two of the 19 breakpoints to lie in the confidence intervals of the S&P 500 due to chance.

Breakpoint 15 is not marked with an asterisk in section 5.1 while the other four breakpoints are marked with an asterisk. Since the Dutch economy is highly dependent on the U.S. economy it is understandable that the words of Greenspan, changes in the federal fund rate, and the Clinton-Lewinsky affair also influence the Dutch AEX. International events such as the Russian default and Asia crisis clearly effect both economies.

6.3 Model

We now come to the basic question how these global findings are related to the generating mechanism. Essentially there are two basic models. The first model is a stochastic model of regime switching based on an underlying Markov chain, the second model is a deterministic dynamic system with several attracting equilibrium points.

There is a growing literature on regime switching in stock markets and in volatility. See Franses and van Dijk (2000) who treat these topics in their recent book, pages 69-205. Regime switching in volatility is treated in terms of GARCH models and an underlying finite state Markov process. In such a model breakpoints occur at random, the regimes have exponentially distributed durations, where the parameter of the exponential distribution may depend on the state.

Often the state space contains only two states. Our results do not support such a two state model. Leaving out the four peaks in the activity level in figure 13, one may try to link each level with a state of a Markov chain. A finite state Markov chain does not seem to fit our data. Setting the state of the middle regime between 1000 and 2000 days at zero, it is not clear what the value of the other regimes should be. For instance, the high plateau around day 3000 may be regarded as state 2 or as state 3.

Since we are looking at a rather long time period and since the ICT revolution is within the time-span of the data, it is possible that the underlying process is Markovian but with a non-stationary state space.

The second model is a deterministic model. Here dynamic laws determine the volatility process. In this setting one can envision a deterministic dynamical system with an attracting fixed point. Unexpected news events with a low intensity and expected news events induce shocks on the level of the state variables. This implies that volatility fluctuates around a mean value. Clearly this model is too simple to account for the bird's eye view of figure 13, where there seem to be many different mean values around which the volatility fluctuates. Indeed, one could extend this model to one with a finite number of fixed points. In such a model unexpected news events with a high intensity may push the volatility into the basin of attraction of a different fixed point. However, this model would have the same flaw as the finite state Markovian model. There is no evidence for a finite state space of fixed points. Moreover, for ten transitions from one equilibrium to another we were not able to find links to news shocks.

A deterministic dynamical model will in general not be time reversible. The system moves towards an equilibrium gradually but leaves the equilibrium by a jump. We have checked for reversibility by applying the distribution free Wilcoxon test. The null hypothesis that the distribution of the deviations from the activity before and after a breakpoint at three and five days are identical can not be rejected.

A sand pile model, see Bak, Tang and Wiesenfeld (1988) with grains of sand replaced by news items might be more appropriate. Most grains that are dropped on the sand pile only change the pile locally, but occasionally a grain may initiate a landslide which changes the whole configuration of the pile.

7 Conclusion

There is evidence of a piecewise linear structure for the quadratic variation of the log returns of the S&P 500 futures index over the period January 1988 until August 2001. This is the period between the October 1987 stock market crash and the September 2001 calamity. The visual evidence of the graph \hat{q} in figure 1 is substantiated by statistical analysis resulting in the piecewise linear approximation in figure 10. Over the 14 year period under observation we distinguish 27 regimes determined by 26 breakpoints.

There is strong evidence of volatility explosions. We see four very thin high peaks in figure 13. As explained above, our procedure for finding breakpoints is not very good for pinpointing such explosions. The four volatility explosions may be linked to news events. Three of these news events do not occur in our list of ten major shocks. So one might argue that the explosions are triggered by news rather than being caused by them.

Our findings do not support a regime switching model for volatility with two states, low volatility and high volatility. There is no clear evidence of a finite state Markov chain in figure 13. As explained above one might argue for a finite state model with changing levels due to non-stationarity.

Approximately sixteen jumps may be linked to news events. These events are of the second order and may be said to precipitate the jumps rather than causing them. Although such jumps are exogenous, the news events which we have recorded may be regarded as a selection from a point process with an intensity in the order of one point per two or three months.

In the alternative model, a smooth dynamical system with several equilibria,

news events may push the system into a basin of attraction of some other equilibrium. If so, there should be signs of time irreversibility, each new regime starting with a gradual approach towards the new equilibrium. The Wilcoxon test does not reject the hypothesis of time reversibility.

Figure 14 plots the activity level in financial time. If we leave out the four volatility explosions, most of the remaining regimes have a length of 0.01 financial time units or a multiple of this base length. This suggests a threshold model and not a Markovian model in which the length of regimes are exponentially distributed.

The behaviour we have observed for the S&P 500 futures index, over the period January 1988, August 2001, also occurs for the Dutch stock index, the AEX, which was observed over the four year period May 1996, September 2000. There is evidence of coincidence of jumps in the two indices.

In this paper we have presented evidence that volatility fluctuates around a fixed level, the activity, over periods of months or even years. Apart from periods of constant activity, we see occasionally volatility explosions. We are not aware of any economic theory which explains this behaviour.

References

- Andersen, T.G. and Bollerslev, T. (1998). Deutsche mark dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance*, **53**, 219–265.
- Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, **66**, 47–78.
- Bai, J. and Perron, P. (2001). Computation and analysis of multiple structural changes. *Forthcoming in Journal of Applied Econometrics*.
- Bak, P. Tang, K. and Wiesenfeld, K. (1988). Self-organized criticality. *Phys. Rev. A*, **38**, 364–374.
- Berry, T.D. and Howe, K.M. (1994). Public information arrival. *Journal of Finance*, **49**, 1331–1446.
- Breiman, L. Friedman, J.H. Olshen, R.A. and Stone, C.J. (1984). *Classification and regression trees*. Belmont: Wadsworth Advanced Books and Software.

- Buffington, J. and Elliott, R.J. (2002). American options with regime switching. *International Journal of Theoretical and Applied Finance*, **5**, 497–514.
- Chong, T.T.L. (2001). Estimating the location and the number of change points by the sample-splitting method. *Statistical Papers*, **42**, 53–79.
- Cutler, D.M., Poterba, J.M. and Summers, L.H. (1989). What moves stock prices? *Journal of Portfolio Management*, **15**, 4–12.
- Dacorogna, M.M, Gençay, R., Müller, U., Olsen, R.B. and Pictet, O.V. (2001). *An Introduction to High-Frequency Finance*. London: Academic Press.
- Davies, P.L. (2003). Approximating long range financial data. Working Paper.
- Diebold, F.X. and Inoue, A. (1999). Long memory and regime switching. *Journal of Econometrics*, **105**, 131–159.
- Duan, J. Popova, I. and Ritchken, P. (2002). Option pricing under regime switching. *Quantitative Finance*, **2**, 116–132.
- Ederington, L.H. and Lee, J.H. (1993). How markets process information: News releases and volatility. *Journal of Finance*, **48**, 1161–1191.
- Fisher, W.D. (1958). On grouping maximum homogeneity. *Journal of the American Statistical Association*, **53**, 789–798.
- Franses, P.H. and van Dijk, D. (2000). *Non-linear Time Series Models in Empirical Finance*. Cambridge: Cambridge University Press.
- Hamilton, J.D. (1989). A new approach to the economic analysis of non-stationary time series and the business cycle. *Econometrica*, **57**, 357–384.
- Jackwerth, J.C. and Rubinstein, M. (1996). Recovering probability distributions from option prices. *Journal of Finance*, **51**, 1611–1631.
- Mikosch, T. and Starica, C. (1999). Change of structure in financial data, long-range dependence and GARCH modelling. Working paper.
- Peters, R.T. and de Vilder, R.G. (2002a). The S&P500 future index as a time changed Brownian motion. Delta Preprint Series 2002-06, Paris.

Peters, R.T. and de Vilder, R.G. (2002b). The realized volatility of the main dutch (AEX) stock index. Appeared as 'I.i.d. standard normality for the Dutch (AEX) stock index' in Delta Preprint Series 2002-05, Paris.

Timmerman, A. (2001). Structural breaks, incomplete information and stock prices. FMG Discussion Papers : dp311.