Physica A 392 (2013) 350–360

Contents lists available at SciVerse ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Modelling financial volatility in the presence of abrupt changes

Gordon J. Ross

Heilbronn Institute for Mathematical Research, University of Bristol, United Kingdom

ARTICLE INFO

Article history:
Received 19 January 2012
Received in revised form 3 August 2012
Available online 10 September 2012

Keywords:
Volatility modeling
GARCH
Change detection
Nonparametric statistics

ABSTRACT

The volatility of financial instruments is rarely constant, and usually varies over time. This creates a phenomenon called volatility clustering, where large price movements on one day are followed by similarly large movements on successive days, creating temporal clusters. The GARCH model, which treats volatility as a drift process, is commonly used to capture this behaviour. However research suggests that volatility is often better described by a structural break model, where the volatility undergoes abrupt jumps in addition to drift. Most efforts to integrate these jumps into the GARCH methodology have resulted in models which are either very computationally demanding, or which make problematic assumptions about the distribution of the instruments, often assuming that they are Gaussian. We present a new approach which uses ideas from nonparametric statistics to identify structural break points without making such distributional assumptions, and then models drift separately within each identified regime. Using our method, we investigate the volatility of several major stock indexes, and find that our approach can potentially give an improved fit compared to more commonly used techniques.

1. Introduction

Volatility clustering is often observed in the return series of financial instruments [1,2]. This phenomena is best illustrated by an example. Let S_t denote the price of some financial instrument at a set of equally spaced discrete time points $t = \{1, 2, \ldots\}$, and let the *return series* be the log-increments $r_t = \log S_t - \log S_{t-1}$. The volatility of the instrument is defined as the standard deviation of these returns. A typical example of a financial return series can be seen in Fig. 1, which shows the daily returns of the Dow Jones stock index over a 20 year period ranging from January 1991 to August 2011. It can be observed that the standard deviation is not constant, but instead varies over time. In particular, note that the period from 2003 to 2007 seems to have noticeably lower volatility than the period immediately before or after. Similarly, in 2008 there are many extreme return values which occur in close succession, pointing to an abnormally high volatility during this period.

Volatility clustering refers to this notion that large/small returns tend to be followed by similarly large/small values, which results in extended regimes of abnormally high or low volatility. This has been empirically observed in many different financial time series, and poses a problem for traditional financial models, which have typically assumed that the volatility is roughly constant over time. The last 25 years have seen an increasing number of attempts to model the time-varying nature of volatility, and the generalised autoregressive conditional heteroskedasticity (GARCH) model [3], along with its many variants, is now the de-facto standard. The idea behind GARCH is that the volatility undergoes a stochastic drift process, where the conditional volatility at time t is a random variable, with a conditional distribution which depends on the long term volatility, the volatility during the most recent period, and the most recent values of the return series.

However the gradual drift process underlying the GARCH model seems to be empirically violated in many real financial series. In some cases, volatility seems to behave more like a jump process, where it fluctuates around some value for an extended period of time, before undergoing an abrupt change, after which it fluctuates around a new value. This can be seen

E-mail addresses: gr203@ic.ac.uk, gordon.ross03@imperial.ac.uk.

0378-4371/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.physa.2012.08.015

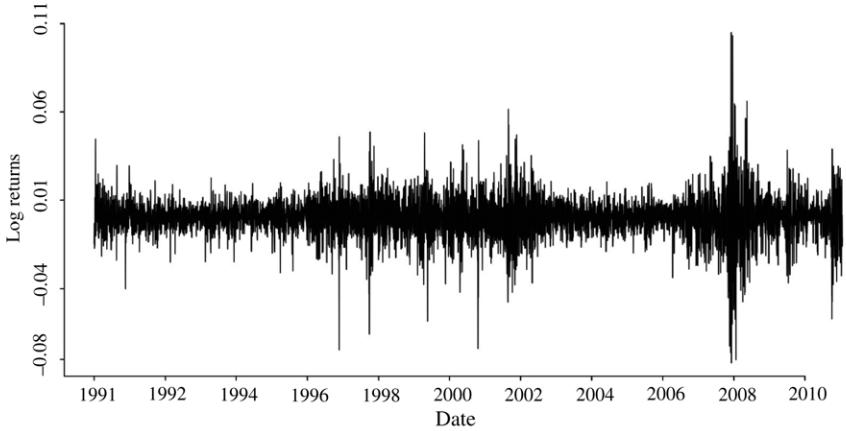

Fig. 1. Daily returns of the Dow Jones index between January 1984 and August 2011.

in Fig. 1 around the year 1996, where the volatility spontaneously increases for a period of several years, before dropping to a lower value during 2003. Since the standard GARCH model does not contain the possibility of these sudden jumps, it tends to overestimate the degree of long term volatility persistence. This has prompted the development of regime-switching GARCH processes which can incorporate jumps [4,5]. In these models, the return series is allowed to contain multiple change points which segments it into regimes, with the GARCH model having different parameters within each segment.

However, such models can be hard to estimate. Although there are computationally efficient procedures for estimating multiple change points in simpler ARCH models [6–8], the long-range dependence introduced by the GARCH formulation makes such approaches difficult to apply. Standard techniques for fitting multiple change point models to data assume independence between segments [9] which is not the case in the GARCH framework. Although some recent attempts to fit such models have been attempted [5], it remains a difficult numerical procedure. Therefore, the most popular strategy is to instead use the approximate procedure introduced by [10] where the model is fitted in stages, with the abrupt change points first being located using the iterated cumulative sum of squares (ICSS) algorithm [11], before a GARCH model then estimated conditional on these change points. This ICSS–GARCH algorithm has been used to study a wide variety of financial time series. For example, [12] uses it to study the volatility of the US dollar exchange rate against several different currencies, [13] studies the returns of the Canadian stock exchange, [14] does likewise for the Japanese and Korean exchanges, and [15] analyses the market for crude oil.

Although ICSS–GARCH is simple to implement and has been shown to give improved results compared to standard GARCH models, it is not without its problems [16,17]. The parameters of the ICSS algorithm are usually designed under the assumption that the financial returns follow a Gaussian distribution, and it can produce many spurious jump points if this assumption is violated. We will show later that applying ICSS to heavy-tailed series can give poor results, since extreme observations are misinterpreted as being regime shifts. Unfortunately, it has now been conclusively established that financial data is very rarely Gaussian, and return series typically exhibit heavy-tail behaviour [18–22]. Similar heavy-tail behaviour has been observed in many financial series and is not limited to asset returns [23].

This limitation of the ICSS–GARCH methodology has meant that it is usually only used to detect change points in the weekly returns of financial instruments, i.e. where S_t and S_{t+1} are one week apart [24,14,15,13,10]. Using the algorithm on daily returns can generate too many spurious false positives for it to be useful, due to the number of extreme values. This is a problem since the daily returns are more fine-grained and hence using them should allow more accurate volatility modelling. Therefore, it is desirable to find a way to use this data when it is available.

In this paper we present an alternative to the ICSS–GARCH algorithm which is better suited for dealing with the heavy tailed, non-Gaussian data which is typical in finance. We replace the ICSS segmentation step of ICSS–GARCH with an alternative technique based on non-parametric statistics, which does not make any assumptions about the true returns distribution. This allows it to entirely avoid the Gaussianity assumption and allows it to be deployed on daily returns. Our approach is based on the nonparametric change point model framework described in Ref. [25,26], and we hence refer to it as NPCPM–GARCH. Using this technique, we analyse several stock indexes for volatility change points, specifically focusing on the Dow Jones Industrial Average, the German DAX, the VIX volatility index, and the Japanese Nikkei 225. We compare our results with those of ICSS–GARCH, and find that our method generally gives a better fit to the data when measured using standard criteria. This suggests that it could be a widely useful tool for modelling volatility in other contexts.

The remainder of the paper proceeds as follows. We begin in Section 2 by describing the ICSS step of the ICSS–GARCH algorithm. We explain why it gives poor performance when used with heavy-tailed data, and give a simulated example using Student-t to show this. In Section 3 we introduce our new nonparametric approach. We then briefly review the GARCH stage of the algorithm in Section 3.3, and in Section 4 we present an empirical evaluation of our method on a range of foreign exchange series.