Journal of Empirical Finance 17 (2010) 485–500

Contents lists available at ScienceDirect

Journal of Empirical Finance

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

Tracking a changing copula

Andrew Harvey

Faculty of Economics, University of Cambridge, United Kingdom

ARTICLE INFO

Article history:
Received 16 September 2008
Received in revised form 22 July 2009
Accepted 16 October 2009
Available online 30 October 2009

JEL classification:
C14
C22
F36

Keywords:
Concordance
Contagion
Exponentially weighted moving average
Quantiles
Signal extraction
Tail dependence

ABSTRACT

A copula models the relationships between variables independently of their marginal distributions. When the variables are time series, the copula may change over time. Recursive procedures based on indicator variables are proposed for tracking these changes over time. Estimation of the unknown parameters is by maximum likelihood. When the marginal distributions change, pre-filtering is necessary before constructing the indicator variables on which the recursions are based. This entails estimating time-varying quantiles and a simple method based on time-varying histograms is proposed. The techniques are applied to the Hong Kong and Korean stock market indices. Some interesting and unexpected movements are detected, particularly after the attack on the Hong Kong dollar in 1997.

1. Introduction

Stock returns are known to be non-normal with a distribution that changes over time. The most pervasive form of time variation is changing variance or volatility. However, features other than scale, such as skewness or kurtosis, may also change.

Just as the normal distribution is inadequate for modeling univariate time series, so the bivariate normal distribution is not suitable for modeling the relationship between two assets. As well as the asset returns not being normally distributed, their comovements may not be adequately captured by correlation coefficients. For example, marginal distributions tend to be characterized by fat tails and the probability of two markets both exhibiting a relatively high movement (in the same direction) may be much higher than can typically be captured with a bivariate normal distribution. A copula models the relationships between two variables independently of their marginal distributions. It does so by means of a joint distribution function with standard uniform marginals. Hence it gives the probability that the observations in two series are below certain quantiles.

There is evidence to suggest that copulas may sometimes change over time; see for example, Van den Goorbergh et al. (2005), Rodriguez (2007) and Patton (2006). Tracking the movements in different parts of the copula may point to a variety of changes in the relationship between the two series. In particular we may wish to focus on movements in dependence as characterized by the probability that one series is below (above) a given quantile, given that the other is below (above) a given quantile. If a single measure of dependence is required, it may be appropriate to consider the probability that both observations are below their respective medians. A simple transformation of this measure yields *Blomqvist's beta*, which, because it lies in the range $[-1, 1]$, is comparable with other measures of association; see Kruskal (1958) and Fermanian and Scaillet (2003).

E-mail address: ach34@econ.cam.ac.uk.

0927-5398/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2009.10.004

The aim of this article is to propose some simple statistical procedures for tracking copula probabilities over time and to investigate how these procedures fare with real data. The distribution function of the copula is not specified, so the approach is different from the one employed by Patton (2006) where the parameters in the copula are assumed to be functions of past observations. Instead probabilities associated with the copula are estimated by using filters designed for binary time series. A straightforward option is to employ an exponentially weighted moving average (EWMA) of binary indicators, defined according to whether observations from two series both lie below particular quantiles. The discount parameter in the EWMA is estimated by a maximum likelihood. Smoothed estimates, based on two-sided filters, can be computed by drawing on the correspondence with the Gaussian local level model.

When the medians are constant over time, Blomqvist's beta is unaffected by changes in volatility. However, even if the medians are constant, which is not necessarily the case for stock returns, other parts of the copula will certainly be affected by changes in the marginal distributions. Hence some kind of pre-filtering is needed. The most general solution is to try to track the distribution functions of the marginals. The approach adopted here is based on generalizing the filter for binary observations so as to deal with categorical data.[1] The categories correspond to parts of the distribution and the filters for the proportions in each category are EWMAs, as in the binary case. The discount parameter, or parameters, can be estimated by maximizing a likelihood function based on the multinomial distribution. Given these proportions, in what may be regarded as a time-varying histogram, quantiles can be estimated at each point in time by interpolation.[2]

The plan of the paper is as follows. Section 2 discusses the method for estimating the changing probability in a binary time series and indicates its relevance to time-varying copulas. Section 3 sets out the proposal for estimating time-varying quantiles and explores the connection between this method and non-parametric procedures, as in Yu and Jones (1998).

Figs. 1 and 2 show daily returns in the Hong Kong (Hang Seng) and Korean (SET) stock price indices[3] from 27/11/79 to 27/11/07. Three events are marked on Fig. 2: (i) Black Monday, October, 19th, 1987; (ii) the speculative attack on Hong Kong dollar on 20 October 1997; and (iii) the 'high tech.' crash of 2nd October 2000. The increase in volatility in Fig. 1 immediately after 20th October 1997 is clearly discernible. (The large negative return was on Black Monday.) Section 4 uses the techniques developed in Sections 2 and 3 to explore the relationship between the two indices, with special emphasis on the issue of contagion stemming from the attack on the Hong Kong dollar; see Dungey et al. (2005). Tracking the copula provides a coherent description and yields some new insights.

Section 5 returns sets out a method, similar to that adopted in Section 3 for computing changing proportions in univariate distributions, for simultaneously estimating copula frequencies defined for a grid on the unit square. Such estimates might be used to compute estimates of changing measures of association, such as Spearman's rank correlation coefficient and Kendall's Tau. Section 6 concludes.

2. Tracking changes in the copula

2.1. Copulas

A copula is a joint distribution function of standard uniform random variables, U_1 and U_2, that is

$$C(u_1, u_2) = \Pr(U_1 \leq u_1, U_2 \leq u_2), \quad 0 \leq u_1, u_2 \leq 1.$$

When the variables are independent, $C(u_1, u_2) = \Pr(U_1 \leq u_1). \Pr(U_2 \leq u_2) = u_1 u_2$.

The copula gives the probability that observations on two variables, Y_1 and Y_2 are less than or equal to given quantiles, that is

$$C(\tau_1, \tau_2) = \Pr(Y_1 \leq \xi_1(\tau_1), Y_2 \leq \xi_2(\tau_2)) = F(\xi_1(\tau_1), \xi_2(\tau_2)), \quad t = 1, \ldots, T, \tag{1}$$

where $\xi_i(\tau_i)$ is the τ_i-quantile for $i = 1, 2$. The probability that both observations lie above their pre-assigned quantiles is known as the survival function and it is equal to

$$\overline{C}(\tau_1, \tau_2) = \Pr(Y_1 > \xi_1(\tau_1), Y_2 > \xi_2(\tau_2)) = 1 - \tau_1 - \tau_2 + C(\tau_1, \tau_2); \tag{2}$$

see, for example, Cherubini et al. (2004, p75) or McNeil et al. (2005, p196). Note that $\bar{C}$ (0.5, 0.5) = C(0.5, 0.5).

The copula provides a flexible way of capturing dependence. The quadrant association, $\bar{C}(\tau_1, \tau_2) + C(\tau_1, \tau_2)$, gives a measure of dependence in the range [0, 1]; see Kruskal (1958, p 818). *Blomqvist's beta*, $2(\bar{C}(0.5, 0.5) + C(0.5, 0.5)) - 1 = 4C(0.5, 0.5) - 1$, lies in the range $[-1, 1]$ and is zero when the series are independent.

[1] An alternative approach would be to estimate a time-varying distribution computed using a kernel. Such an approach is currently under investigation.

[2] Of course, if the researcher is prepared to assume a parametric model, such as GARCH or stochastic volatility, better estimates of the time-varying quantiles may result, provided that the model is well-specified.

[3] The Hong Kong and Korean stock price indices are based on local currency – the Hong Kong dollar and Korean won respectively.