Chapter 1

Introduction

The aim of this thesis is to examine the application of parameter-mixing to multivariate statistical models, and in particular the properties of copula-based statistical models constructed or extended through parameter-mixing. Copula functions offer a convenient way via which flexible, multivariate statistical models can be constructed. The question of interest here is how parameter-mixing, a method frequently used to introduce new models in statistical analysis, can be usefully applied to copula functions.

Traditionally, when studying multivariate distributions, the focus has been on the multivariate normal distribution; see, for example, Anderson [1958]. Multivariate normal distributions have the appealing property that every marginal distribution is also normal. However, in many instances the multivariate normal distribution is not appropriate to the situation. For example, Joe [1997] lists longitudinal count, binary, and ordinal response data as examples of non-normal variables. In such cases, a non-normal multivariate distribution is required; see, for example, Johnson and Kotz [1972]. However, in modelling, the specification of non-normal multivariate distributions can often be problematic for several reasons. Firstly, each family of multivariate distributions usually corresponds to a specific set of joint distributions; changing the specification of marginal distributions requires re-specifying a different family of joint distributions. Secondly, while it is relatively simple to extend univariate distributions to the two-dimensional case (such as the bivariate Poisson distribution), it is more difficult to further extend them to higher dimensions. Finally, it is generally difficult to describe the dependence structure independently of the margins (Frees and Valdez [1998]).

These problems can be alleviated through the specification of multivariate models via copulas. The copula is a function that binds together univariate distributions to form a multivariate distribution. It captures the dependence structure between the variables, separate from the margins. Hence, a multivariate distribution can be specified by choosing any set of univariate distributions together with a copula function. Because every multivariate distribution can be expressed in copula form (Sklar's theorem; see below Chapter 2), the copula provides a useful way to specify families of multivariate models using any desired marginal distributions.

Families of copulas are often indexed by a parameter. Since the copula captures the precise dependence structure among random variables, this parameter is termed a dependence parameter, with its value determining the dependence structure captured by the particular copula. With given data, dependence parameters can be estimated using standard methods such as maximum likelihood estimation.

Although the use of copulas began relatively recently in econometrics, there is a growing body of literature on their application. Cherubini et al [2004] details the application of copulas in finance. In a time series context, a series of work by Patton (e.g. [2006]) demonstrate the application of copulas, and in particular, copulas conditional on past information, to modelling exchange rates. In the microeconometric context, Smith [2003] demonstrates the application of copulas to model selectivity bias without relying on multivariate normality.

As discussed above, a central motivation for using copula-based models is to avoid the rigidity associated with multivariate normal and other multivariate distributions. It is worthwhile, then, to investigate whether methods can be found to further enhance the flexibility of a copula-based model. One common method by which statistical models are made more flexible is through (what is termed here) parameter-mixing. In general, a parameter-mixture distribution arises from a hierarchical model. Starting from a *parent* distribution for the variable or variables of interest, its parameter (or set of parameters) is assigned a *mixing distribution*. This achieves a hierarchical model, where a distribution depends on a parameter which, in turn, is described by another distribution. The density function for the *mixture distribution* can then be found by taking expectations with respect to the original parameter. Parameter-mixing is discussed in detail in Johnson et al [1993]. If successfully applied, parameter-mixing can generate a new model which has improved modelling properties over the original. One example is the *Beta-Binomial* distribution. In this case, the univariate $Binomial(n, p)$ distribution is generalised by assuming that the parameter p (success probability) has a $Beta(\alpha, \beta)$ distribution. The resultant univariate $Beta\text{-}Binomial(n, \alpha, \beta)$ distribution allows greater flexibility by successfully extending a two-parameter distribution to a three-parameter one, where each of the parameters is separately identified.

The application of parameter-mixing to copulas has been discussed on several occasions. Mikusiński et al [1991] discuss with some detail the construction and probabilistic interpretations for mixture copulas. Nelsen [2006] also discusses the process of generating new copulas via the parameter-mixing process, and shows that the mixed function is also a copula. However, the existing literature neglects the link between the mixture copula and its parent copula. Questions which have not been addressed include whether the relationship between mixture copulas and their parents can be exploited advantageously, and whether, and when, a mixture copula displays modelling advantage over its parent copula. In addition, the existing literature on parameter-mixture copulas focus almost exclusively on one particular family: Uniform mixtures of shuffles of C (see, for example, Ferguson [1995]). The modelling properties of other mixture copulas has not been discussed. This thesis aims to address these issues by investigating parameter-mixing as applied to a number of copulas commonly used in modelling, with the aim of ascertaining whether parameter-mixing results in modelling advantages, and the circumstances under which they arise.

First, the theoretical background to copula-based modelling and the method of parameter-mixing, especially its application to copulas, is reviewed in Chapter 2. Chapter 3 demonstrates the construction of a copula by parameter-mixing and investigates some properties of parameter mixture copulas. In particular, the link between the mixture copula and the parent copula is exploited to yield some advantageous results.

In order to investigate the modelling properties of the mixture copula, an experiment is conducted in Chapter 4 using simulated data. By comparing the modelling results of the mixture copula as compared to the parent copula, the experiment demonstrates the circumstances under which mixture yields modelling advantage, and the extent of this advantage. Finally, mixture copula estimation is applied to a real life data set in Chapter 5. In particular, the use of parameter-mixing to convey prior information is investigated.