Chapter 2

Copulas and mixture copulas

2.1 Copulas

The copula for a p-dimensional multivariate distribution function H, given its marginal distribution functions $F_1(x_1)$, $F_2(x_2)$, ... , $F_p(x_p)$, binds these marginal distributions together so as to generate the joint distribution function H. That is, the copula is the function C such that

$$H(x_1, x_2, ...x_p) = C(F_1(x_1), F_2(x_2), ..., F_p(x_p)) \tag{2.1}$$

where $F_i(x_i)$ is the cumulative density function (cdf) of x_i such that $F_i(x_i) = \Pr(X_i \leq x_i)$; by definition of a cdf, $0 \leq F_i(x_i) \leq 1$. As the copula function can be specified separately to the marginal distributions, the choice of marginal distributions is not constrained.

One interpretation of (2.1) is that a multivariate distribution function is composed of two parts: marginal distribution functions for each of the variables, and a multivariate structure, embodied in the copula (Joe [1997]). Alternatively, we can view the copula as a function that "binds" the marginal distributions together to form the multivariate distribution. While the copula can have arbitrary dimensions, most of its properties can be sufficiently illustrated in two dimensions, and extend readily to higher dimensions. Hence, the following discussion will focus on two-dimensional or bivariate copulas, denoted as 2-copulas.

Nelsen [2006] gives a definition of a 2-copula, which can be derived from (2.1), and using the properties of a bivariate distribution. Let the unit interval be $\mathbb{I} = [0, 1]$, and the unit square $\mathbb{I}^2$ be the product $\mathbb{I} \times \mathbb{I}$. Under Nelsen's definition, a 2-copula is a function C from $\mathbb{I}^2$ to $\mathbb{I}$, such that:

1. For every u, v in $\mathbb{I}$,

$$C(u, 0) = 0 = C(0, v)$$

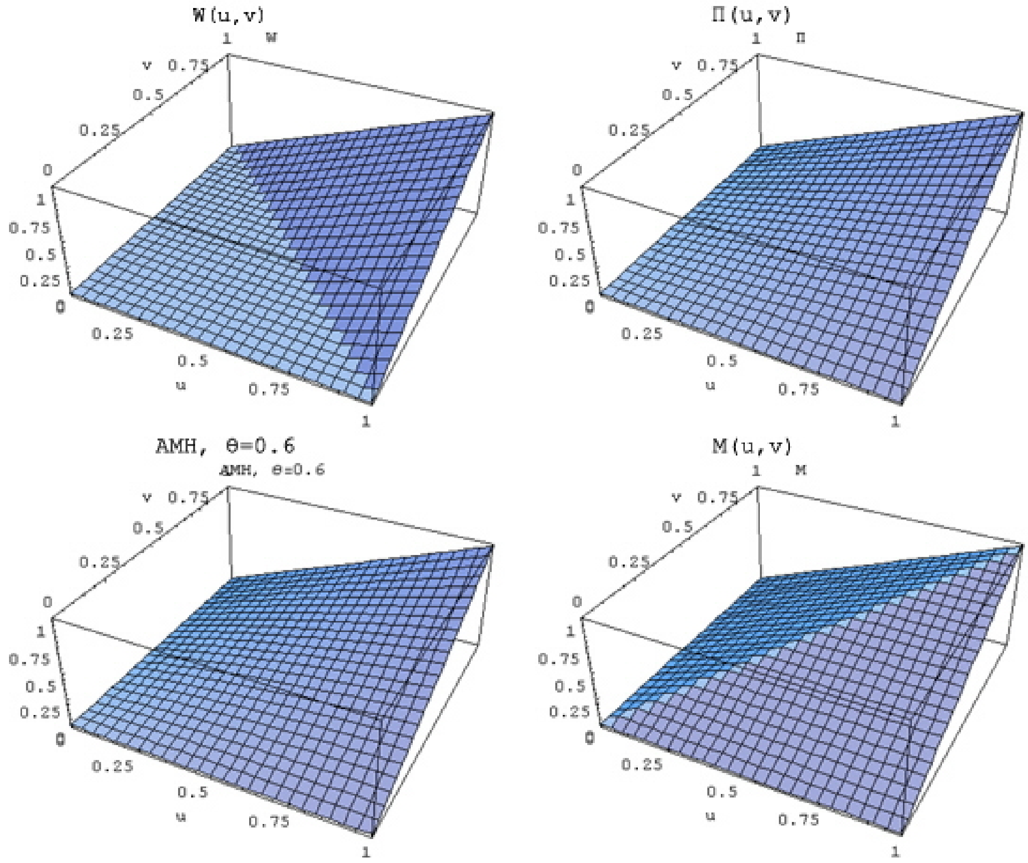

Figure 2.1: Plots of some copulas

and

$$C(u, 1) = u \text{ and } C(1, v) = v$$

2. For every u_1, u_2, v_1, v_2 in $\mathbb{I}$ such that $u_1 \leq u_2$ and $v_1 \leq v_2$,

$$C(u_2, v_2) - C(u_2, v_1) - C(u_1, v_2) + C(u_1, v_1) \geq 0. \tag{2.2}$$

These features are illustrated in Figure 2.1.

The name "copula" was first coined by Sklar [1959]. However, study of these functions predate the use of the term; see, for example, Hoeffding [1940], and Fréchet [1951]. Kimeldorf and Sampson [1975] use the term "uniform representation", and Galambos [1978] uses the term "dependence functions". Copulas are also similar to t-norms found in the study of probabilistic metric spaces in mathematics; see, for example, Schweizer and Sklar [1983, Section 6.3]. For a comprehensive study of the origin and development of copulas, see Schweizer [1991], or Sklar [1996].

2.2 Sklar's theorem

A crucial theorem that underlies most applications of the copula was proposed by Sklar [1959]. From (2.1), it is readily apparent that, given any marginal distributions $F_1(x_1)$, $F_2(x_2)$, ... , $F_p(x_p)$, and a copula C, we can construct the joint multivariate distribution H. Sklar's theorem establishes the converse: given H, the joint distribution function with margins $F_1(x_1)$, $F_2(x_2)$, ... , $F_p(x_p)$, there exists a copula C such that (2.1) holds. If the marginal distributions are continuous, then C is unique. Otherwise, if the marginal distributions take on only a non-continuous lattice of values, then C is uniquely determined on these values. Schweizer and Sklar [1974] demonstrate one proof of this theorem in two dimensions; the n-dimensional proof is given by Schweizer and Sklar [1983, Theorem 6.2.5], and also by Moore and Spruill [1975, Lemma 3.2].

The importance of Sklar's theorem lies in its main implication, that any joint distribution can be expressed in copula form. This means that, for any set of random variables characterised by some joint distribution function, there exists a unique copula C (or unique on the range of values of the respective marginal distributions). Hence, any problem involving multiple random variables can be couched in terms of a copula, so that the copula can be applied to many situations involving consideration of a multivariate distribution.

Sklar's theorem has several further implications. Firstly, this theorem allows us to construct copulas from a given joint distribution and its margins. Taking a bivariate version of (2.1), we have:

$$H(x_1, x_2) = C(F_1(x_1), F_2(x_2)) \tag{2.3}$$

As demonstrated by Nelsen [2006], from this we can obtain:

$$C(u, v) = H(F_1^{(-1)}(u), F_2^{(-1)}(v)) \tag{2.4}$$

where $F_1^{(-1)}$ and $F_2^{(-1)}$ are quasi-inverse functions[1] of F_1 and F_2, respectively.

Furthermore, since independence between a pair of random variables implies that their joint distribution is the product of the marginal distributions, it follows from Sklar's theorem that the copula that corresponds to the independence case must be the "product copula", $\Pi(u, v) = uv$.

[1] A quasi-inverse $F^{(-1)}$ of a function F is such that, if t is within the range of F, then

$$F(F^{(-1)}(t)) = t;$$

otherwise, $F^{(-1)}(t)$ takes the nearest boundary value on the domain of F; that is:

$$F^{(-1)}(t) = \inf\{x|F(x) \geq t\} = \sup\{x|F(x) \leq t\}.$$

If F is strictly increasing, then the quasi-inverse is the ordinary inverse, F^{-1}.

2.3 Families of Copulas

Many copulas can be organised into "families" indexed by parameters. For modelling purposes, parameterised copulas have the advantage that they are able to be applied to given data in order to investigate the correlation structure among the random variables of interest. Let θ be a parameter (possibly vector valued). The parameterised 2-copula can be written as:

$$C_\theta(u, v)$$

Since the copula captures the dependence structure in the joint distribution, the value of the parameter θ determines the association in the corresponding bivariate distribution. Hence it is called the *dependence parameter*. Provided that the margins F_1 and F_2 do not depend on θ, (2.3) holds for any $C_\theta(u, v)$.

A large variety of copula families exist. One example is the Farlie-Gumbel-Morgenstern (FGM hereafter) family of copulas, which has the form:

$$C_\theta(u, v) = uv + \theta uv(1-u)(1-v) \qquad \text{where } -1 \leq \theta \leq 1. \tag{2.5}$$

Because of its simple form, the FGM copula can be used to construct simple and analytically tractable distributions. It has been used in modelling, tests of association, and in studying the efficiency of nonparametric procedures (Hutchinson and Lai [1990]).

Some families of copulas fall within the class known as "Archimedean copulas", which are characterised by a generator function, $\varphi : \mathbb{I} \to [0, \infty]$, where φ is a continuous and strictly decreasing function. Any such φ then generates a valid bivariate cdf from its margins. In copula form,

$$\varphi(C(u, v)) = \varphi(u) + \varphi(v) \tag{2.6}$$

or:

$$C(u, v) = \varphi^{(-1)}(\varphi(u) + \varphi(v)) \tag{2.7}$$

where $\varphi^{(-1)}$ is the quasi-inverse defined above. Where the generator depends on a dependence parameter θ, it can be expressed as φ_θ, in which case a whole family of copulas is Archimedean and is indexed by θ.

The Ali-Mikhail-Haq (AMH hereafter) is an example of an Archimedean family of copulas. The AMH family is characterised by the generator function

$$\varphi_\theta(t) = \begin{cases} \ln \frac{1-\theta(1-t)}{t}, & -1 \leq \theta < 1 \\ \frac{1}{t} - 1, & \theta = 1. \end{cases}$$

Since $\varphi_\theta(0) = \infty$, this generator is said to be *strict*, and $\varphi_\theta^{(-1)} = \varphi_\theta^{-1}$, by the definition of the quasi-inverse. Hence, the quasi-inverse of this function is, for $\theta \in [0, 1)$:

$$\begin{aligned}\varphi_\theta^{(-1)}(t) &= \varphi_\theta^{-1}(t)\\ &= \frac{1-\theta}{e^t - \theta}\end{aligned}$$

and for $\theta = 1$:

$$\begin{aligned}\varphi_1^{(-1)}(t) &= \varphi_1^{-1}(t)\\ &= \frac{1}{1+t}.\end{aligned}$$

Using (2.7), the AMH family of copulas is then given by, for $\theta \in [0, 1)$:

$$\begin{aligned}C_\theta(u,v) &= \frac{1-\theta}{\exp[\ln\frac{1-\theta(1-u)}{u} + \ln\frac{1-\theta(1-v)}{v}] - \theta}\\ &= \frac{(1-\theta)uv}{[1-\theta(1-u)][1-\theta(1-v)] - \theta uv}\\ &= \frac{(1-\theta)uv}{1-\theta(2-u-v+uv)+\theta^2(1-u)(1-v)}\\ &= \frac{(1-\theta)uv}{(1-\theta)-\theta(1-\theta)(1-u)(1-v)}\\ &= \frac{uv}{1-\theta(1-u)(1-v)}\end{aligned}$$

and for $\theta = 1$:

$$\begin{aligned}C_1(u,v) &= \frac{1}{1+(\frac{1}{u}-1)+(\frac{1}{v}-1)}\\ &= \frac{uv}{u+v-uv}\\ &= \frac{uv}{1-(1-u)(1-v)}.\end{aligned}$$

Hence, the AMH family of copulas is given by:

$$C_\theta(u,v) = \frac{uv}{1-\theta(1-u)(1-v)}, \qquad -1 \le \theta \le 1. \tag{2.8}$$

Another example is the Gumbel-Barnett (GB) copula, which is characterised by the generator function $\varphi_\theta(t) = \ln(1-\theta\ln t)$. Since $\varphi_\theta(0) = \infty$, this generator is also strict, and the quasi-inverse of the generator is

$$\begin{aligned}\varphi_\theta^{(-1)}(t) &= \varphi_\theta^{-1}(t)\\ &= \exp\left[\frac{1-e^t}{\theta}\right].\end{aligned}$$

Hence the GB family of copulas is given by:

$$\begin{aligned} C_\theta(u,v) &= \exp\left[\frac{1-\exp[\ln(1-\theta\ln u)+\ln(1-\theta\ln v)]}{\theta}\right] \qquad \text{where } 0<\theta\le 1 \\ &= \exp\left[\frac{1-(1-\theta\ln u)(1-\theta\ln v)}{\theta}\right] \\ &= \exp\left[\frac{\theta(\ln u+\ln v)-\theta^2\ln u\ln v}{\theta}\right] \\ &= uv\exp\left(-\theta\ln u\ln v\right). \end{aligned} \tag{2.9}$$

2.4 Dependence Coverage

An important property of a family of copulas is its dependence coverage. The ranges of values that θ can take defines the range of dependence structures which a family of copulas is capable of capturing. There are three special cases of copulas which are relevant to discussions of dependence coverage, illustrated in Figure 2.1. The first is the Product copula, usually denoted by Π:

$$\Pi(u,v) = uv \tag{2.10}$$

which represents independence of the margins.

The other two are denoted W and M respectively:

$$W(u,v) = \max(u+v-1,0) \tag{2.11}$$

$$M(u,v) = \min(u,v). \tag{2.12}$$

These two copulas form the Fréchet-Hoeffding lower bound and Fréchet-Hoeffding upper bound, respectively. In the bivariate case, W represents perfect negative correlation; M represents perfect positive correlation. These copulas are the universal upper and lower bounds of any bivariate copula, such that, for any copula C and all u, v in $\mathbb{I}$,

$$W \le C(u,v) \le M.$$

Because of Sklar's theorem, these actually represent the bounds for any bivariate joint distribution function H.

Within these bounds, each family of copulas usually occupies a subset of permissible copula forms, representing a subset of the entire range of permissible dependence structures. The range of dependence structures covered by the family is called its "dependence coverage", and the usefulness of a copula family in modelling can often depend on its dependence coverage.

A number of statistics can be used to measure and compare the dependence structures captured by particular copulas. The most commonly used measure of dependence is Pearson's correlation coefficient. However, Pearson's correlation coefficient is not invariant with respect to margins, and is not a good measure of non-linear correlation. By contrast, "scale-invariant" dependence measures have the advantage of being invariant to margins (Joe [1997]). Two most frequently used scale-invariant dependence measures are the Spearman Rank Correlation Coefficient (Spearman's rho) and the Kendall Rank Correlation Coefficient (Kendall's tau).

Happily, any scale-invariant dependence measure can be expressed in terms of the copula. This follows from a result shown by Schweizer and Wolff [1981], that the copula corresponding to the distribution of a set of random variables is invariant under strictly increasing transformations of the random variables. More formally, let X_1 and X_2 be continuous random variables with copula C_{X_1,X_2}, as in (2.3). If functions α and β are strictly increasing functions defined on the ranges of X_1 and X_2 respectively, the copula of $\alpha(X_1)$ and $\beta(X_2)$ is the same as that of X_1 and X_2; that is $C_{\alpha(X_1),\beta(X_2)} = C_{X_1,X_2}$. In other words, the copula is invariant under strictly increasing transformations of X_1 and X_2. This result shows that, under increasing transformations of the random variables, the copula is invariant while the margins can be changed at will. Thus, the copula captures those aspects of the joint distribution that are invariant to increasing transformations of the random variables As a result, any dependence measure that remains unchanged under strictly increasing transformations of the random variables (that is, it is "scale invariant") can be expressed in terms of the copula. Expressing Spearman's rho and Kendall's tau (denoted here by ρ and τ respectively, or ρ_θ and τ_θ respectively in the case of a family of copulas) in copula terms, we have

$$\rho = 12 \iint_{\mathbb{I}^2} C(u,v)dudv - 3 \tag{2.13}$$

and

$$\begin{aligned} \tau &= 4 \iint_{\mathbb{I}^2} C(u,v)dC(u,v) - 1 \\ &= 4E[C(U,V)] - 1 \end{aligned} \tag{2.14}$$

respectively. For parametric copulas, these formulas can often reduce to a relatively simple expression in terms of the dependence parameter.

For example, the Spearman's rho and Kendall's tau measure of the AMH copula (2.8) can be expressed as follows (Nelsen [2006, Section 2.4]):

$$\rho_\theta = \frac{12(1+\theta)}{\theta^2}\mathrm{di}\log(1-\theta) - \frac{24(1-\theta)}{\theta^2}\ln(1-\theta) - \frac{3(\theta+12)}{\theta} \tag{2.15}$$

$$\tau_\theta = \frac{3\theta-2}{3\theta} - \frac{2(1-\theta)^2}{3\theta^2}\ln(1-\theta).$$

By evaluating these expressions at various values of the parameter θ, the dependence coverage can be found. For the AMH family, dependence coverage measured in Spearman's rho is

$$\rho_\theta \in [33 - 48\ln 2, 4\pi^2 - 39] \cong [-0.2711, 0.4784].$$

When $\theta = 0$, the AMH copula becomes the product copula.

A small dependence coverage can limit the usefulness of a copula family. For example, the FGM copula has dependence coverage $\rho_\theta \in [-\frac{1}{3}, \frac{1}{3}]$, which, as Joe [1997] notes, restricts the usefulness of this family for modelling.

2.5 Modelling using copulas

The copula is a versatile modelling tool that allows highly flexible multivariate models to be constructed. By separating the specification of marginal distributions from the specification of the correlation structure, multivariate models can be constructed with arbitrary marginal specification. This avoids the problems commonly faced by conventional multivariate model construction, such as the need to specify a different family for each combination of marginal distributions, and the difficulty in comparing dependence structures between models with different marginal distributions.

Underlying copula-based modelling is Sklar's theorem, which states that any multivariate distribution can be specified as the marginal distributions linked by a copula function. Thus, building a statistical model using copulas involves a two-step process. Firstly, marginal distributions for each variables are specified. In a parametric context, selection of the marginal distributions can be conducted with standard estimation procedures. Secondly, a copula is chosen that binds the marginal distributions together to form the multivariate distribution.

In the bivariate case, consider two random variables X_1, X_2. The goal is to estimate a model for the true but unknown joint distribution of X_1 and X_2. To construct this model, propose two marginal distributions, $F_1(x_1)$, $F_2(x_2)$, and a parameterised copula family, C_θ. Together, this specifies a model for the joint cdf for both variables.

The greatest advantage of the copula method is that choice of marginal distributions is unconstrained. Since the copula is unique upon the range of values of the margins (by Sklar's theorem), margins need not be identical to each other, and furthermore they can be continuous or discrete, or a mixture of both. Additionally, since the copula captures the association between the variables and is invariant to changes in the margins, the copula parameters comprehensively describe the dependence structure. Also, as the copula is invariant to the margins, location and scale can be specified via specification of the margins, freeing the copula parameters to measure the dependence structure.

The application of copula-based models in economics and econometrics is an expanding area of research. Copula-based models have been applied to a range of multivariate problems. For example, Joe [1997] illustrates several examples of modelling using copulas. More recently, Smith [2003] uses Archimedean copulas to construct sample selection models; Cameron et al [2004] apply copulas to model count data, and specifically count data measured with error; Smith [2005] applies copulas to model wage earnings models for child workers. 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. The extensive application of copulas in finance is detailed by Cherubini et al [2004].

Copula-based models can be estimated by maximum likelihood, generally with a two-stage process. Two methods are commonly used. The first, parametric method is termed Inference Functions For Margins (IFM), which requires parameterised marginal distributions to be specified. Each margin is in turn estimated via maximum likelihood, and the fitted margins are then used to estimate the copula via maximum likelihood. The second, semi-parametric method is termed Canonical Maximum Likelihood (CML), where, instead of the fitted margins, the empirical cdf is used to estimate the copula. Estimation by IFM is illustrated in Chapter 6.

2.6 Parameter-mixing

Since the central motivation for using copulas is to avoid rigidity in the choice of margins and specification of correlation structure, it is worthwhile to examine methods for enhancing the flexibility of copula-based models. One method which is frequently used to enhance the flexibility of parametric models is parameter-mixing.

In parametric statistics, the parameter is usually assumed to be fixed. Relaxing this assumption serves to introduce added flexibility to a model. A parameter-mixture distribution (hereafter mixture distribution) arises from a hierarchical model, where the quantity or quantities of interest are assigned a distribution, which depends on a parameter that also has a distribution itself. The hierarchical model can be used to model a complicated process by placing successive, simple, models in a hierarchy. Since the quantity or quantities of interest follow a distribution that depends on a quantity that also has a distribution, it is said to follow a mixture distribution.

More formally, let X be a random variable whose *parent distribution* $\mathcal{F}_A$ with cdf $F_\theta(x|\theta_1, \theta_2, ..., \theta_m)$ depends on a set of parameters $\theta_1, \theta_2, ..., \theta_m$, and some or all of these parameters vary according to some *mixing distribution* $\mathcal{F}_B$. Then, X is said to follow a *mixture distribution*.

The mixture distribution can be understood to involve $m + 1$ variables, being X,

$\Theta_1, \Theta_2, ...\Theta_m$, with the parent distribution corresponding to the conditional distribution of $X|(\Theta_1 = \theta_1, \Theta_2 = \theta_2, ...\Theta_m = \theta_m)$.

The mixture distribution function is simply the marginal distribution function of X that can be obtained by taking the expectation of the parent distribution with respect to the mixing distribution:

$$\begin{aligned} F(x; \lambda_1, ..., \lambda_p) &= E_{\mathcal{F}_B}[F_\theta(x|\Theta_1, \Theta_2, ...\Theta_m)] \\ &= \int_{\Theta_1} ... \int_{\Theta_m} F_\theta(x|\theta_1, \theta_2, ..., \theta_m) dF_B(\theta_1, \theta_2, ..., \theta_m; \lambda_1, ..., \lambda_p) \end{aligned} \tag{2.16}$$

where the notation indicates that the mixing distribution may depend upon a set of p parameters $\lambda = \{\lambda_1, ..., \lambda_p\}$ which originate from the mixing distribution $\mathcal{F}_B$ that has cdf $F_B(\theta_1, \theta_2, ..., \theta_m; \lambda_1, ..., \lambda_p)$. A notation in common use to denote parameter-mixing is

$$\mathcal{F}_A \underset{\Theta}{\wedge} \mathcal{F}_B$$

where $\mathcal{F}_A$ is the parent distribution and $\mathcal{F}_B$ is the mixing distribution.

To extend an existing model, the parameter of interest is often assigned a distribution that depends on a parameter set larger in dimension than the base, or parent, distribution; in other words, $p \geq m$. A famous example of this being successfully achieved is the *Beta-Binomial* distribution, which can be represented as:

$$Binomial(n, P) \underset{P}{\wedge} Beta(\alpha, \beta).$$

Starting from the 2-parameter univariate $Binomial(n, p)$ distribution, the success probability p is assigned a $Beta(\alpha, \beta)$ distribution. Put formally, this parameter-mix is the following expectation with respect to the parameter distribution:

$$\begin{aligned} Beta\text{-}Binomial(n, \alpha, \beta) &= Binomial(n, P) \underset{P}{\wedge} Beta(\alpha, \beta) \\ &= E[Binomial(n, P)]. \end{aligned}$$

The $Binomial(n, p)$ distribution is interpreted as a conditional distribution, conditioning on the value p of the variable P. The resultant $Beta\text{-}Binomial(n, \alpha, \beta)$ distribution is well known because it successfully extends a two-parameter model to a three-parameter model, each of whose parameters are formally identified.

2.7 Parameter-mixing applied to copulas

Given the past success of models constructed through parameter-mixing, it is interesting to investigate whether the technique, applied to copulas, could improve the properties of copula models. To establish notation, for a (possibly vector-valued) dependence parameter $\Theta = \theta$ assigned the distribution $\mathcal{F}(\lambda)$, where λ are themselves parameters assumed to have dimensionality at least that of Θ, the $\mathcal{F}$ parameter-mix of the 2-copula C_θ is given by

$$\begin{aligned} C'_\lambda(u,v) &= C_\Theta(u,v) \underset{\Theta}{\wedge} \mathcal{F}(\lambda) \\ &= E\left[C_\Theta(u,v)\right] \\ &= \int_\Theta C_\theta(u,v) dF(\theta;\lambda) \\ &= \int_\Theta C_\theta(u,v) f(\theta;\lambda) d\theta \end{aligned} \tag{2.17}$$

where $F(\theta;\lambda)$ denotes the cdf of Θ and $f(\theta;\lambda)$ the pdf (probability density function) of Θ; the last line assumes continuity of the mixing distribution. The outcome $C'_\lambda(u,v)$ (hereafter termed a mixture copula) is a 2-copula. It is clear that, having averaged over the values of Θ, the range of dependence structures covered by C'_λ cannot exceed that covered by C_θ. In fact, it can be less if $\mathcal{F}(\lambda)$ is used to represent prior information.

There has been some detailed discussion on mixture copulas as a method of constructing new copulas. Another term by which it is known is the *convex sum* of C_θ with respect to $\mathcal{F}$; see, for example, Nelsen [2006, Sec. 3.2.4]. This name arises from an alternative interpretation of parameter mixing. Firstly, it is easy to verify that any convex linear combination of a finite collection of copulas is a copula satisfying the definition in (2.2). The result can then be extended to an infinite collection of copulas C_θ indexed by a parameter θ. Taking θ as an observation on a continuous random variable Θ with distribution $\mathcal{F}(\lambda)$, the convex sum of this family of copulas, weighted by the density of Θ is then the mixture copula, as in (2.17).

Mikusiński et al. [1991] discuss the construction and probabilistic interpretation of the mixture copula. In particular, they discuss a particular family of copulas, known as "shuffles of C", which has the form:

$$C_\theta(u,v) = \begin{cases} C(1-\theta+u,v) - C(1-\theta,v), & u \leq \theta \\ v - C(1-\theta,v) + C(u-\theta,v), & u > \theta \end{cases} \tag{2.18}$$

They show that, with a $Uniform(0,1)$ mixing distribution, the resultant mixture copula is the product copula Π.

Discussions on *modelling* using mixture copulas has been relatively scant. Extant literature on the subject focus almost exclusively on mixture models based on shuffles of particular copulas with a uniform mixing distribution, as in (2.18). One example is Ferguson [1995], who discusses a model that can be interpreted as the $Uniform(0,1)$-mixture of a copula derived by "shuffling". There has been no investigation of modelling using mixture copulas other than Uniform mixtures of shuffles of C.

Thus, the mechanism by which copulas can be constructed by mixing, and the interpretations thereof, have been discussed in the past. However, the literature on modelling using mixture copulas is small, and, where it exists, focuses almost exclusively on models based on Uniform mixtures of shuffles of C. As a result, the modelling properties of mixture copulas has received little attention. In particular, the link between mixed and unmixed copulas (mixture copulas and their parents) has not been explored. In examining the issue of modelling with mixture copulas, this thesis focuses on the link between mixed and unmixed copulas. In particular, this means examining whether and when mixture copulas have modelling advantages over their unmixed counterparts. In other words, whether and when mixing results in modelling advantage.