Copulas

In statistics, a copula can be thought of as defining the correlation structure for a family of multivariate distributions. Each distribution in the family is constructed by "gluing" together two or more univariate distributions, with the copula supplying the "glue". Copulas can be applied to a wide range of simulation problems - for example, in financial modelling.

At Mark 23 of the NAG Fortran Library, in addition to the Gaussian and Student's t copulas, we have available the copulas:

- Clayton/Cook-Johnson bivariate (g05ref);
- Frank bivariate (g05rff);
- Plackett (g05rgf);
- Clayton/Cook-Johnson (g05rhf);
- Frank (g05rjf);
- Gumbel-Hougaard (g05rkf).

Bivariate copulas use the conditional sampling approach as different from the less computationally efficient mixture of powers method adopted by the general multivariate copulas.

To illustrate, suppose we wish to simulate the joint distribution of two or more random variables. The kind of copula we use - Frank, Student's t, etc. - describes the correlation structure between the variables, while a collection of univariate distributions (usually called the marginal distributions) define the distribution within each of the variables. Compared to more conventional multivariate distributions (such as, for example, the multivariate Normal), the use of a copula allows each variable to have a different marginal distribution. Thus, for example, one variable may have a Gaussian marginal distribution, while another has a Beta distribution, a third is uniformly distributed, and so on.

The below graphs display a bivariate distribution from a Frank copula with Beta marginal distributions. The input for the calculation of the bivariate distribution includes its copula parameter (usually θ), which contains details of the correlation between the different variables. The copula is created using NAG function g05rff from a repeatable sequence initialised by g05kff. The shape of the Beta distribution is defined by two parameters (usually α and β), and can range from the symmetric bell curve of the Normal distribution (when α and β are both large) to being highly skewed or, at the other extreme when α and β are both 1, a uniform distribution. Here we set $\alpha = 12$ and $\beta = 5$ for the first variable - which gives a skewed distribution – and $\alpha = \beta = 5$ for the second variable, which is distributed symmetrically.

10,000 Beta bivariates with Frank copula dependence with θ=3

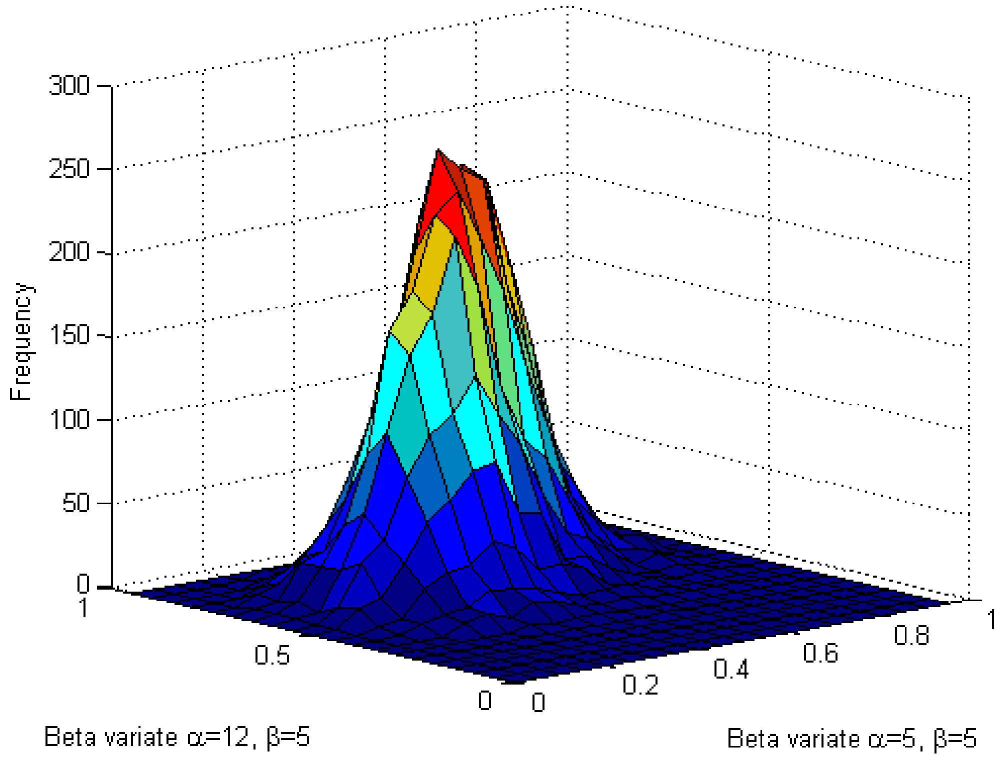

10,000 Beta bivariates with Frank copula dependence with θ=3