Copulas and Dependence

MFM Practitioner Module: Quantitative Risk Management

John Dodson

January 24, 2018

Review: Random Variables on $\mathbb{R}^+$

The most important parametric random variable with half-line support is the Generalized Inverse Gaussian

Generalized Inverse Gaussian (GIG)

$$f(x) = \frac{\chi^{-\lambda}\left(\sqrt{\chi\psi}\right)^{\lambda}}{2K_{\lambda}\left(\sqrt{\chi\psi}\right)} x^{\lambda-1} e^{-\frac{\chi}{2x}-\frac{\psi x}{2}}$$

for $x > 0$, where $K_{\lambda}(\cdot)$ is modified Bessel function of the second kind.

- This generalizes the Gamma and reciprocal Gamma
- There are several versions of parameterization in use
- Other members of this family include the inverse Gaussian and the reciprocal inverse Gaussian
- The name relates to the first passage time of a Brownian motion through a boundary

Mixtures

We saw last semester that conditioning generally reduces entropy. Mixing has the opposite effect. This is useful if you want to moderate statistical hubris.

Say X is the parametric r.v. we are interested in.

1. Concatenate it with some function of the parameters to make a multivariate r.v. $X \# \Theta$.
2. Specify the marginal density of Θ and the conditional density of $X|\Theta$.
3. Integrate over the support of Θ to get the marginal density of X, the new mixture.

$$
\begin{aligned}
& f_{X \# \Theta}(x \# \theta) = f_{X|\Theta}(x;\theta) f_{\Theta}(\theta) \\
\implies & f_X(x) = \int_{\Omega(\Theta)} f_{X|\Theta}(x;\theta) f_{\Theta}(\theta)\, d\theta
\end{aligned}
$$

Mixtures

Student's-t is a symmetric r.v. which exhibits leptokurtosis.

(Gosset's) Student's-t

Consider a normal r.v. with an unknown variance close to one. If the variance is a draw from an reciprocal Gamma,

$$X \left| \sigma^2 \right. \sim \mathcal{N}\left(0, \sigma^2\right)$$
$$\sigma^2 \sim \text{Gamma}^{-1}\left(\frac{\nu}{2}, \frac{\nu}{2}\right)$$

the resulting unconditioned density is

$$f_X(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)\Gamma\left(\frac{1}{2}\right)} \frac{1}{\sqrt{\nu}} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$$

- The version with $\nu = 1$ is the Cauchy
- The limit $\nu \to \infty$ is a normal

Mixtures

The version of the Student's-t above has a variance for $\nu > 2$, but it is not unity.

Standardized Student's-t

The standardized version can be useful for fitting residuals*. It has the density

$$f_X(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)\Gamma\left(\frac{1}{2}\right)} \frac{1}{\sqrt{\nu-2}} \left(1 + \frac{x^2}{\nu-2}\right)^{-\frac{\nu+1}{2}}$$

- *Note that $\mathrm{E}\, e^X \to \infty$ for any finite ν so Student's-t cannot be used with log-returns of asset prices.
- For historical reasons, if the parameter ν is an integer, it is termed the degrees of freedom.

Mixtures

The generalized hyperbolic family is a normal mean / GIG variance mixture. The Student's-t is a special case (with $\lambda = -\nu/2$).

Normal / reciprocal inverse Gaussian (NRIG)

Another useful GH is the symmetric Normal / reciprocal inverse Gaussian mixture (with $\lambda = \frac{1}{2}$). The standardized version has the density

$$f_X(x) = \frac{1}{\pi} e^g \sqrt{1+g}\, K_0\left(\sqrt{g^2 + (1+g)x^2}\right)$$

for shape parameter $g \geq 0$. It is not obvious, but the limit $g \to \infty$ corresponds to the normal.

- As a model for the residuals of the log-returns of asset prices, this is superior to the Student's-t example from the text because $\operatorname{E} e^X$ is finite.

Mixtures

We can use the NRIG to illustrate the effects of scaling and mixing on entropy.

$$H_X = \overset{\text{(mixing)}}{\mathrm{E}\, H_{X|\Theta}(\Theta)} + H_\Theta = \overset{\text{(scaling)}}{H_{aX}} - \log|a|$$

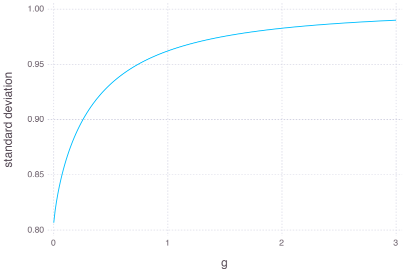

NRIG relative to a Gaussian for fixed entropy

- The leptokurtosis of the NRIG allows the standard deviation to be up to almost 20% lower

Multivariate Extensions

Many (univariate) distributions on the real line can be generalized to (multivariate) random variables on a vector space

- the normal
- Cauchy, Student's-t
- symmetric generalized hyperbolic

Often the replacement

$$\left(\frac{X-\mu}{\sigma}\right)^2 \to (X-\mu)'\Sigma^{-1}(X-\mu)$$

and normalization is all that is required to lift the sample space from $\mathbb{R}$ to $\mathbb{R}^d$.

- This is a rich source for elliptical random variables

Copulas

Parametric multivariate random variables involve at least one class of univariate random variable in the form of the marginals for the components. But it is also clear that the characterization of the original multivariate r.v. is not simply a collection of these marginal characterizations. There is a structure, with its own parameters, that connects them together.

- This is the copula.

For me, this is the prototypical example; and multivariate random variables are a rich source for parametric copulas. But it is not the only source. In fact, any random variable whose sample space is a unit hypercube with standard uniform margins is a copula.

Sklar's Theorem

To the extent that the *joint* density is not just a product of the *marginal* densities, there is dependence.

Factorization

This ratio can be expressed as

$$c\left(F_{X_1}(x_1), F_{X_2}(x_2), \ldots\right) \triangleq \frac{f_{(X_1,X_2,\ldots)}(x_1, x_2, \ldots)}{f_{X_1}(x_1) f_{X_2}(x_2) \cdots}$$

Copula

Sklar's theorem says this is always possible. More generally, $c = f_U : [0,1]^d \mapsto \mathbb{R}^+$ is a density function that characterizes a new random variable, U, that encapsulates the dependence structure of X.

Note that independence means $c \equiv 1$

Normal (Gaussian) Copula

When dependence can be entirely described by correlation, the Gaussian copula can be appropriate. For $d = 2$,

$$c(u) = \frac{1}{\sqrt{1-\rho^2}} \exp\left[\frac{-\rho}{1-\rho^2}\left(\rho\,\mathrm{erfc}^{-1}(2u_1)^2 \cdots \right.\right.$$
$$\left.\left. + \rho\,\mathrm{erfc}^{-1}(2u_2)^2 - 2\,\mathrm{erfc}^{-1}(2u_1)\,\mathrm{erfc}^{-1}(2u_2)\right)\right]$$

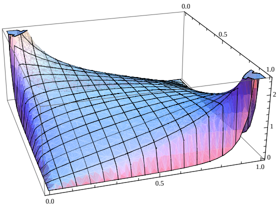

Gaussian copula density for $\rho = \frac{1}{2}$

Tail Dependence

Upper & Lower Tail Dependence

Tail dependence is a pair-wise measure of the concordance of extreme outcomes.

$$\lambda_U = \lim_{p\uparrow 1} \mathrm{P}\left\{X > F_X^{\leftarrow}(p) | Y > F_Y^{\leftarrow}(p)\right\}$$

$$\lambda_L = \lim_{p\downarrow 0} \mathrm{P}\left\{X \leq F_X^{\leftarrow}(p) | Y \leq F_Y^{\leftarrow}(p)\right\}$$

The normal copula fails to exhibit tail dependence: extreme outcomes are essentially independent.

This is a problem, because in practice an extreme outcome in one component often acts to cause extreme outcomes in other components. Developing practical alternatives that include this contagion effect is an active area of research.

Measures of Concordance

Several measures of concordance have been developed. Their definitions are motivated by the properties of their estimators, which we will not discuss just yet. Each ranges from -1 to 1, with 0 for independence. In order of generality, we have

1. Pearson's rho. This is the classical linear correlation measure $\operatorname{cov}(X,Y)/\sqrt{\operatorname{var} X \operatorname{var} Y}$.
2. Spearman's rho. This is linear correlation applied to the grades, $F_X(X)$. It is a simple measure of dependence that is not sensitive to margins.
3. Kendall's tau. This is based strictly on the rank order of pairs of observations of pairs of components. It has useful theoretical and statistical properties.

N.B.: While independence implies zero concordance (under any of these definitions), zero concordance does not imply independence.

Kendall's tau

Kendall's tau can be defined as

$$\tau = 4\,\mathrm{E}\,C(U_1, U_2) - 1$$

where C is the distribution function characterizing the copula of X. It is the probability of concordance minus the probability of discordance for two independent draws of X.

Relationship with other measures

In general Spearman's rho is bounded by

$$\frac{3|\tau| - 1}{2}\operatorname{sgn}\tau \quad \& \quad \frac{1 + 2|\tau| - \tau^2}{2}\operatorname{sgn}\tau$$

For a Gaussian copula, Pearson's rho is

$$\rho = \sin\left(\tfrac{\pi}{2}\tau\right)$$

We use this to define pseudo-correlation for any elliptical r.v..

Normal Mixture Copulas

A normal mixture copula is simply the copula from a normal mixture multivariate random variable. The elliptical copula is an important subclass.

Elliptical Copula

An elliptical random variable is described by a mean vector, a dispersion matrix, and a characteristic generator function. It should be clear that the mean vector has no role in the copula. It should also be clear that the diagonal entries of the dispersion matrix also do not play a role.

Generally, an elliptical copula is parameterized by a semi-definite matrix with unit diagonals, which describe pair-wise dependence, and one or several shape parameters related to the characteristic generator.
The Gaussian copula is an example. Another important example is the t_ν copula, which we will work with in this week's exercise.

Archimedean Copulas

There are on the order of $d^2/2$ parameters to estimate for an elliptical copula. If you are dealing with a very large dimension, such as in a retail or securitization context, you either need a factor model to reduce the dimension or you should consider an Archimedean copula.

Archimedean Copulas

An Archimedean copula is defined in terms of a generator, a decreasing continuous function $\psi : [0, \infty) \mapsto [0, 1]$ with $\psi(0) = 1$ and $\lim_{t\to\infty} \psi(t) = 0$. The copula distribution is

$$C(u_1, u_2, \ldots, u_d) = \psi\left(\psi^{-1}(u_1) + \psi^{-1}(u_2) + \cdots + \psi^{-1}(u_d)\right)$$

Three common single-parameter examples are the Gumbel, Clayton, and Frank.