The Trc. normal distribution

The Truncated normal distribution is a distribution concentrated on the set of positive numbers. This distribution arises as follows. Let Y be a variable having a Normal distribution with some suitable expectation and standard deviation. If we condition on the event that Y is positive, the resulting conditional distribution for Y, is a Truncated normal distribution. We refer to this conditional distribution as the truncated version of the corresponding unconditional distribution; in this case, the Normal distribution.

To construct a variable X with a Truncated normal distribution, we sample values from the Normal distribution until a positive value occurs.

In principle there is no limit on the number of times we need to sample before we get a positive value. However, assuming that we sample from a Normal distribution which has most of its probability mass in the positive region, we usually do not have to sample many times before a usable value occurs.

The Truncated normal distribution has a shape that is similar to a Normal distribution, except that the left tail is cut at zero. On the positive part of the axis the density function is equal to the Normal density function multiplied by a suitable “normalizing” constant.

In the Truncated normal distribution the key numbers, “a”, “b” and “c” are interpreted as follows:

“a”
=
The 10%-fractile of Y

“b”
=
The 50%-fractile of Y

“c”
=
The 90%-fractile of Y

where Y is a variable with the Normal distribution which we want the truncated version of.

To get a sensible distribution, the specified values must satisfy:

“a” < “b” < “c”

DynRisk will adjust the numbers further to make the fractiles fit the fractiles of a Normal distribution.

The Normal distribution is a very common statistical distribution. It is used to model uncertainty in many different situations. The tails of the distribution ranges over the entire set of real numbers.

Sometimes, however, it is necessary to restrict the range of the distribution to the set of positive reals. In such cases one may use either the Truncated Normal distribution (or the Censored Normal distribution). In the Truncated Normal distribution negative values are simply “rejected”.

In most cases the location of the distribution will be such that the probability of having a negative value replaced by zero is small. Thus, in particular, one will typically have:

“a” > 0.

You may of course use the Truncated Normal distribution in situations where one or more of the key numbers are negative as well. In such cases the probability of getting a rejected value may be large. This implies that it may take a while before we get a usable value.