Università Commerciale
Luigi Bocconi

Copulas in Risk Management

Prof. Massimo Guidolin

20541 – Advanced Quantitative Methods for Asset Pricing and Structuring

Winter/Spring 2020

Plan of the Lecture

- The drawbacks of the Normal distribution in multivariate modelling
- Threshold correlations as a detection tool and target to reproduce
- Multivariate Normal threshold correlations
- The copula approach
- Normal copula
- t-Student copula

The Drawbacks of the Normal in Multivariate Modelling

Similarly to what is typically found in univariate applications, also the multivariate Gaussian density seems lacking

- It is well known that both the univariate and the multivariate normal distributions provide poor descriptions of financial returns
  - E.g., this is reflected by non-zero skewness and excess kurtosis
  - This finding is robust to modelling conditional densities of GARCH-type
- The normal distribution is convenient but underestimates the probability of large negative returns
- The multivariate normal distribution has similar problems as it underestimates the joint probability of simultaneous large negative returns across assets (market crashes)
- Risk management models built on multivariate normal distribution are likely to exaggerate the benefits of portfolio diversification
  - Think about slicing and tranching CDOs during the real estate bubble
- Idea: build multivariate shock distributions that are not necessarily normal, $z_t \sim D(0, \Upsilon_t)$

Threshold Correlations as a Detection Tool

Threshold correlations are correlations for selected subsets of data

- Here z_t is now a vector of asset specific shocks, $z_{i,t} \equiv z_{i,t}/\sigma_{i,t}$ and where Υ_t is the dynamic correlation matrix

- Need to develop a diagnostic tool useful to quantify the inadequacy of the multivariate normal when it comes to correlation modelling
- Just as standard univariate analysis uses QQ plots to visualize non-normality, bivariate threshold correlations are useful
- Threshold correlations (TCs) are conventional correlations but computed only on a selected subset of the data
- Consider a probability p and define the corresponding empirical percentile for asset 1 to be $r_1(p)$ and similarly for asset 2, $r_2(p)$
  - These empirical percentiles, or thresholds, can be viewed as the unconditional VaR for each asset
- The threshold correlation for probability level p is now defined by:

$$\rho(r_{1,t}, r_{2,t}; p) = \begin{cases} Corr\left(r_{1,t}, r_{2,t} | r_{1,t} \leq r_1(p) \text{ and } r_{2,t} \leq r_2(p)\right) & \text{if } p \leq 0.5 \\ Corr\left(r_{1,t}, r_{2,t} | r_{1,t} > r_1(p) \text{ and } r_{2,t} > r_2(p)\right) & \text{if } p > 0.5 \end{cases}$$

Threshold Correlations as a Target

- We are computing the correlation conditional on both return series being below their p*th* percentile if $p < 0.5$ and above p*th* percentile if $p > 0.5$
- Threshold correlations are informative about the dependence across asset returns conditional on both returns being either large and negative or large and positive
- They tell us about the tail shape of the bivariate distribution
  - When p gets close to 0 or 1 we run out of observations
  - Clearly the most extreme threshold correlations are quite variable and so should perhaps be ignored
  - An interesting pattern: the tc's get smaller when we observe large negative stock and bond returns simultaneously
  - Large positive stock and bond returns have much higher correlation

Figure 9.1 Threshold correlation for S&P 500 versus 10-year treasury bond returns.

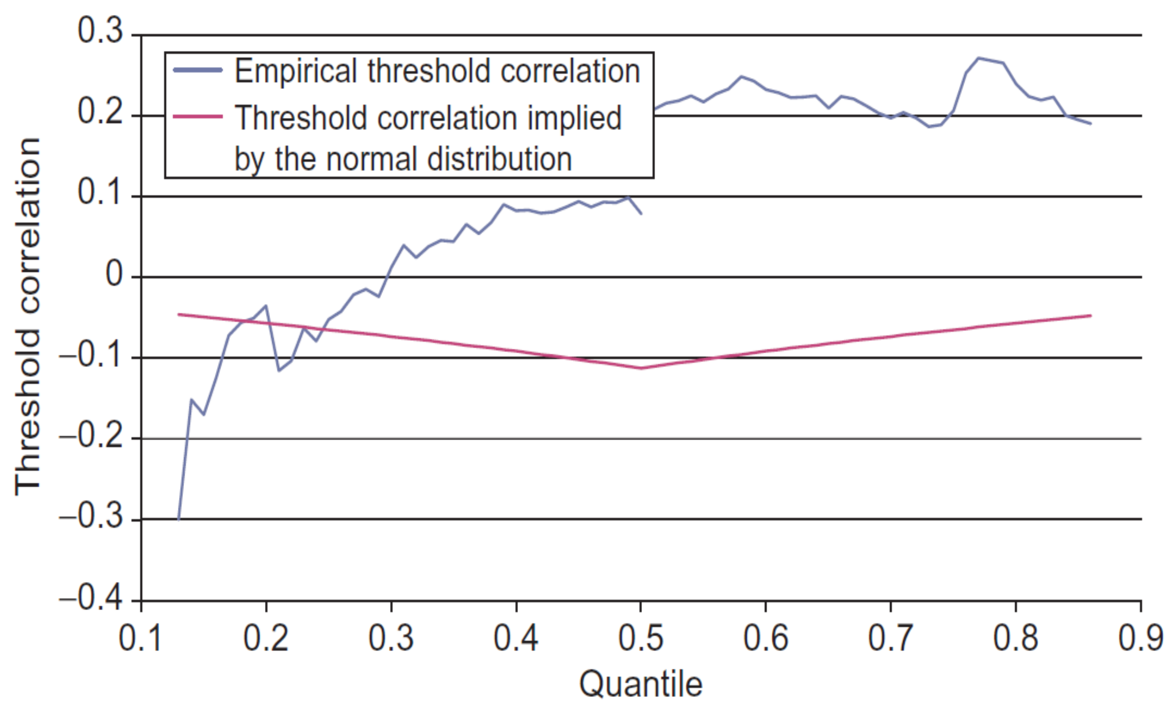

Threshold Correlations as a Target

Important deviations of sample threshold correlations from a multivariate Gaussian indicate the need for better models

- The red line shows the threshold correlations implied by the bivariate normal distribution when using the average linear correlation coefficient

- The bivariate distribution between stock and bonds is asymmetric
- Clearly the normal distribution does not match the threshold correlations found in the data

$$\rho(z_{1,t}, z_{2,t}; p) = \begin{cases} Corr\left(z_{1,t}, z_{2,t} | z_{1,t} \leq z_1(p) \text{ and } z_{2,t} \leq z_2(p)\right) & \text{if } p \leq 0.5 \\ Corr\left(z_{1,t}, z_{2,t} | z_{1,t} > z_1(p) \text{ and } z_{2,t} > z_2(p)\right) & \text{if } p > 0.5 \end{cases}$$

- Given that we are interested in constructing distributions for the return shocks, rather than the returns themselves we compute threshold correlations for standardized shocks, e.g., a GARCH

Figure 9.2 Threshold correlation for S&P 500 versus 10-year treasury bond GARCH shocks.

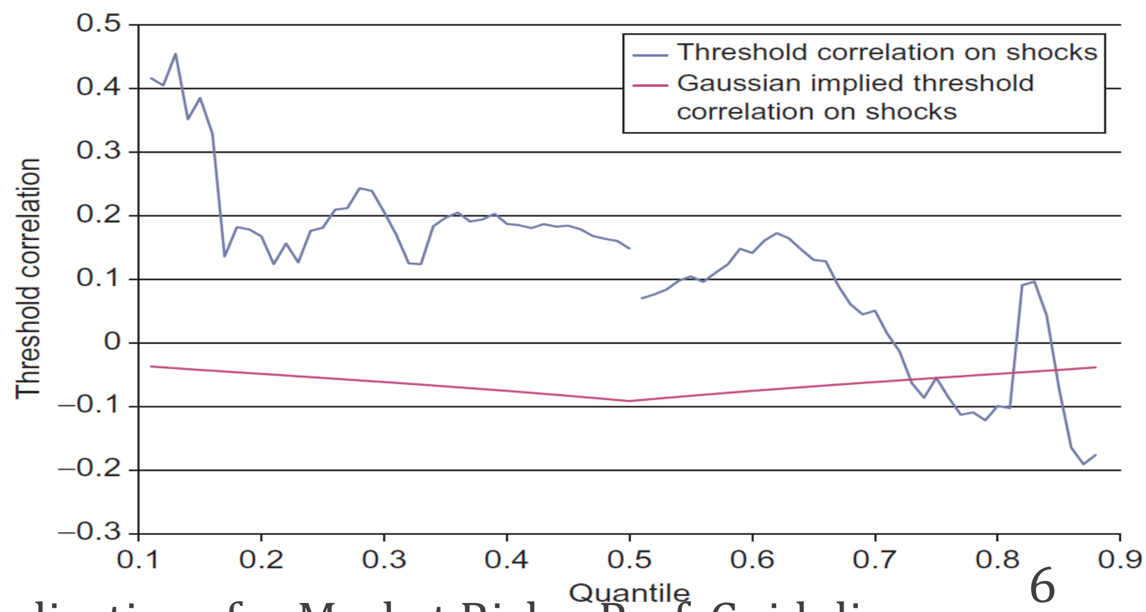

Multivariate Normal Threshold Correlations

Under a multivariate normal benchmark, TCs always go to zero as the quantile goes to 0/1, in the extreme tails variables are uncorrelated

- In any event, stocks and bonds have important nonlinear left-tail dependencies that risk managers need to model

- It is then natural to replace the multivariate normal distribution, i.e., to consider alternative distributions that can be combined with GARCH (or realized variance) and DCC models
- Before we proceed, consider again the multivariate normal benchmark, the standard normal density with correlation ρ:

$$f(z_{1,t}, z_{2,t}; \rho) = \Phi_\rho(z_{1,t}, z_{2,t}) =$$

$$= \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{z_{1,t}^2 + z_{2,t}^2 - 2\rho z_{1,t} z_{2,t}}{2(1-\rho^2)}\right)$$

- Regardless of ρ, TCs go to zero as p → 0 or 1

Figure 9.3 Simulated threshold correlations from bivariate normal distributions with various linear correlations.

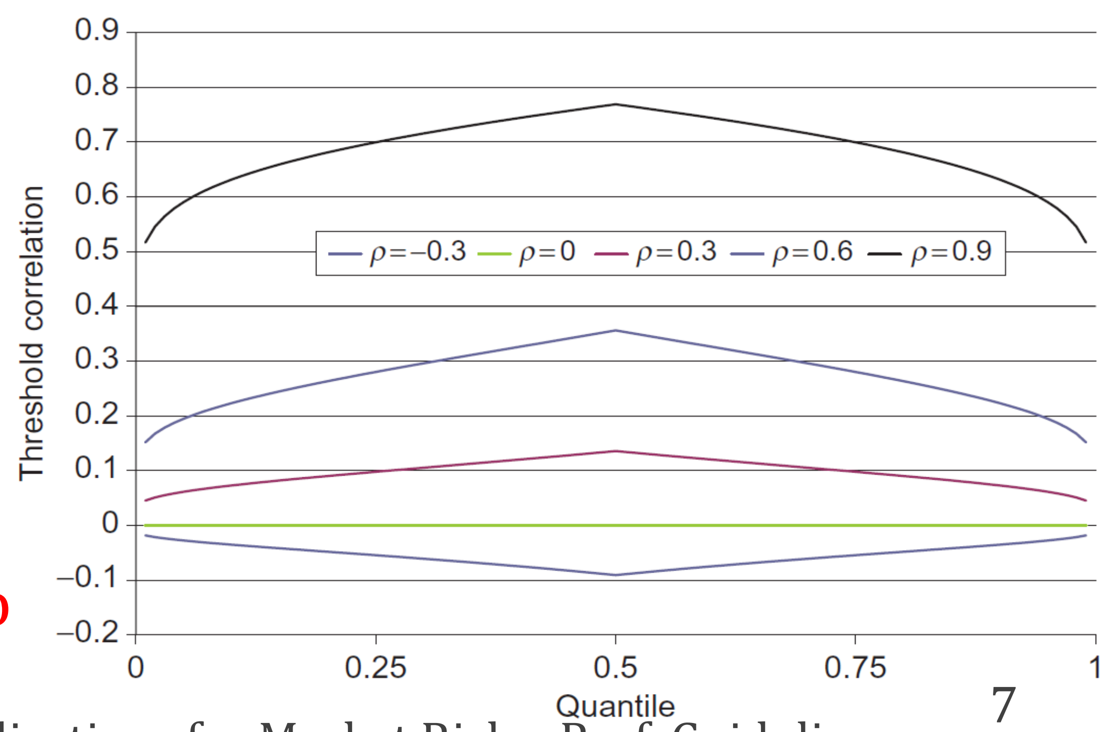

Multivariate t-Student Threshold Correlations

We can generate quite flexible symmetric shapes of tail dependence between the two variables by using a multivariate t distribution

- One popular idea in empirical finance is to replace normal densities with t-Student densities
- Univariate ($d > 2$): $f_{\tilde{t}(d)}(z;d) = \dfrac{\Gamma((d+1)/2)}{\Gamma(d/2)\sqrt{(d-2)\pi}}(1+z^2/(d-2))^{-(1+d)/2}$
- Bivariate (with correlation ρ and $d > 2$):

$$f_{\tilde{t}(d,\rho)}(z_1, z_2; d, \rho) = \frac{\Gamma((d+2)/2)}{\Gamma(d/2)(d-2)\pi\left(1-\rho^2\right)^{1/2}}\left(1+\frac{z_1^2+z_2^2-2\rho z_1 z_2}{(d-2)(1-\rho^2)}\right)^{-(d+2)/2}$$

  - Because d is a scalar, the two variables have the same tail thickness
  - However, we are constrained in one important sense: TCs will always be symmetric vs. the vertical axis

Figure 9.4 Simulated threshold correlations from the symmetric *t* distribution with various parameters.

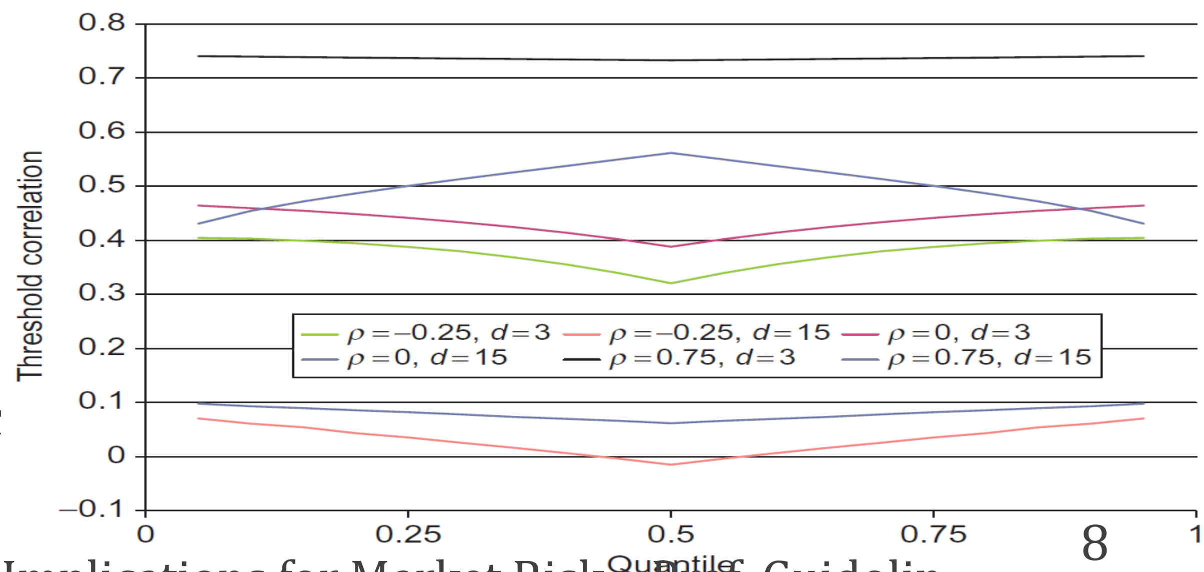

The Copula Approach

Copulas allow us to link the marginal distributions across asset returns to generate a valid multivariate density

- The asymmetric t distribution is able to capture asymmetries in TCs and gaps in the TC around the median (the 0.5 quantile), but hard to use

- Ideally we would like to have an approach where univariate models can be combined to form a proper multivariate distribution
- This is what so-called copula functions have been developed to do
- Consider *n* assets with potentially different univariate (also known as marginal) distributions, $f_i(z_i)$ and cumulative density functions (CDFs) $u_i = F_i(z_i)$ for $i = 1, 2, \ldots, n$
- Copulas help us link the marginals across the assets to generate a valid multivariate density and their use is based on Sklar's Theorem
- For a general class of multivariate CDFs, defined as $F(z_1, \ldots, z_n)$, with marginals $F(z_1), \ldots, F(z_n)$, there exists a unique copula function, G(•) linking the marginals to form the joint distribution

$$F(z_1, \ldots, z_n) = G(F_1(z_1), \ldots, F_n(z_n)) = G(u_1, \ldots, u_n)$$

The Copula Approach

- What's the advantage of copulas? Pay attention:

$$f(z_1,\ldots,z_n) = \frac{\partial^n G(F_1(z_1),\ldots,F_n(z_n))}{\partial z_1 \cdots \partial z_n} = \frac{\partial^n G(u_1,\ldots,u_n)}{\partial u_1 \cdots \partial u_n} \times \prod_{i=1}^{n} f_i(z_i)$$

$$= g(u_1,\ldots,u_n) \times \prod_{i=1}^{n} f_i(z_i)$$

Copula PDF

  - This derives from u_i = integral of $[f_i(z_i) \bullet \partial z_i]$
  - Therefore $\partial u_i = f_i(z_i) \bullet \partial z_i \Rightarrow \partial z_i = [\partial u_i / f_i(z_i)] \Rightarrow (1/\partial z_i) = f_i(z_i)/\partial u_i$

- Consider now the logarithm of the PDF:

$$\ln f(z_1,\ldots,z_n) = \ln g(u_1,\ldots,u_n) + \sum_{i=1}^{n} \ln f_i(z_i)$$

- This decomposition shows that we can build the large and complex multivariate density in a number of much easier steps

  ❶ Build and estimate *n* potentially different marginal distribution models $f(z_1)$, $f(z_2)$, ..., $f(z_n)$ using standard methods

  ❷ Decide on the copula PDF $g(u_1, \ldots, u_n)$ and estimate it using the probability outputs u_i from the marginals as the data

The Copula Approach

Log-likelihood estimation of copula models is made straightforward by the fact that the process may be split in two different steps

- Notice the analogy with GARCH and DCC model building

- This makes high-dimensional modeling possible
  - We can for example allow for each asset to follow different univariate asymmetric t distribution each estimated one at a time
- Sklar's theorem is general – holds for a large class of distributions
- However it is not very specific: It does not say anything about the functional form of G(•) and thus g(•)
- In order to implement the copula modelling approach we need to make specific modeling choices for the copula CDF
- Two choices, among many, are as simple as widespread: the normal and the t-Student copulas
- Most convenient copula is the standard normal, in bivariate case

$$G(u_1, u_2; \rho^*) = \Phi_{\rho^*}(\Phi^{-1}(u_1), \Phi^{-1}(u_2)) = \Phi_{\rho^*}(\Phi^{-1}(F_1(z_1)), \Phi^{-1}(F_2(z_2)))$$

The Normal Copula

The normal copula allows for the marginals to be nonnormal, which in turn can generate many alternative nonnormal distributions

where ρ^* is the correlation between $\Phi^{-1}(u_1)$ and $\Phi^{-1}(u_2)$ and we will refer to it as the copula correlation

- If the two marginal densities, F_1 and F_2, are standard normal then

$$G(u_1, u_2; \rho^*) = \Phi_{\rho^*}(\Phi^{-1}(\Phi(z_1)), \Phi^{-1}(\Phi(z_2))) = \Phi_{\rho^*}(z_1, z_2) \quad \text{(biv. normal)}$$

  - If the marginal distributions are NOT the normal then the normal copula does NOT imply the normal distribution

- The normal copula is more flexible than the normal distribution as the normal copula allows for the marginals to be non-normal, which in turn can generate a multitude of non-normal distributions

Figure 9.6 Simulated threshold correlations from the bivariate normal copula with various copula correlations.

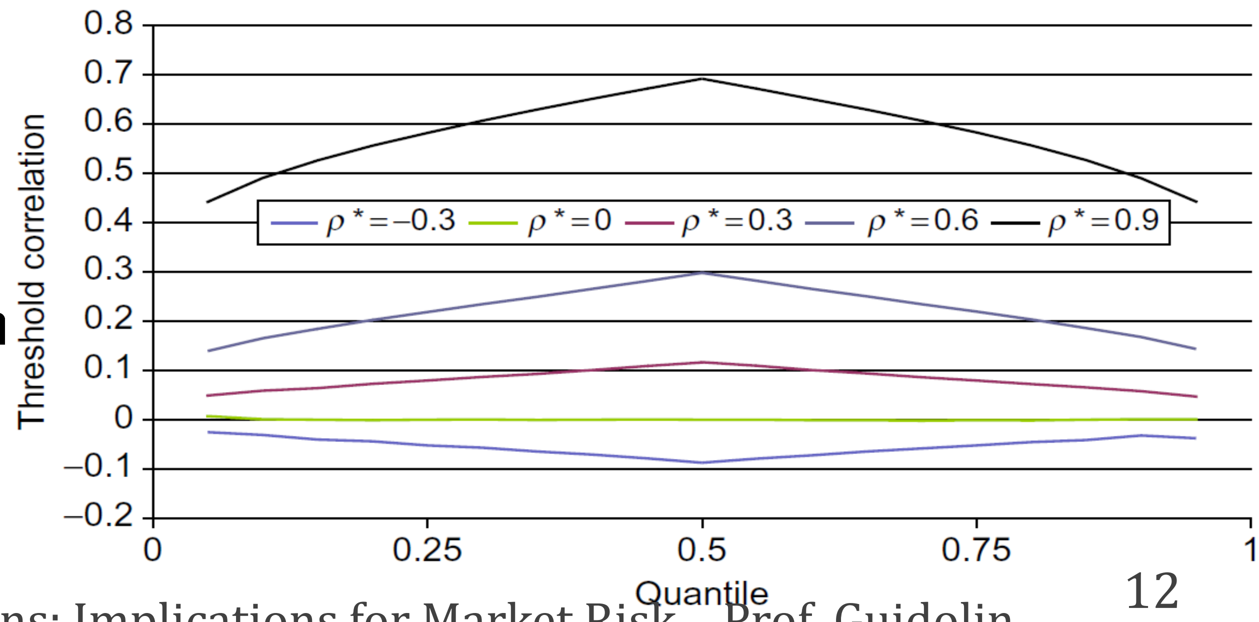

The t-Student Copula

The normal copula cannot generate asymmetric or discontinuous threshold correlations; hence t-Student copulas are needed

- Naturally, the normal copula threshold correlations look similar to the normal distribution threshold correlations
  - The normal copula gives us flexibility by allowing the marginal distributions F_1 and F_2 to be flexible
  - Yet the multivariate aspects of the normal distribution remains: The threshold correlations go to zero for extreme u_1 and u_2, which is likely undesirable when extreme moves are highly correlated across assets
- Fortunately a copula model can be built from the t distribution
- The bivariate t copula CDF is defined by

$$G(u_1, u_2; \rho^*, d) = t_{(d,\rho^*)}\left(t^{-1}(u_1; d), t^{-1}(u_2; d)\right)$$

- Naturally, the t copula threshold correlations look similar to the t distribution threshold correlations
- The t copula can generate large threshold correlations for extreme moves in the assets

The t-Student Copula

Figure 9.7 Simulated threshold correlations from the symmetric *t* copula with various parameters.

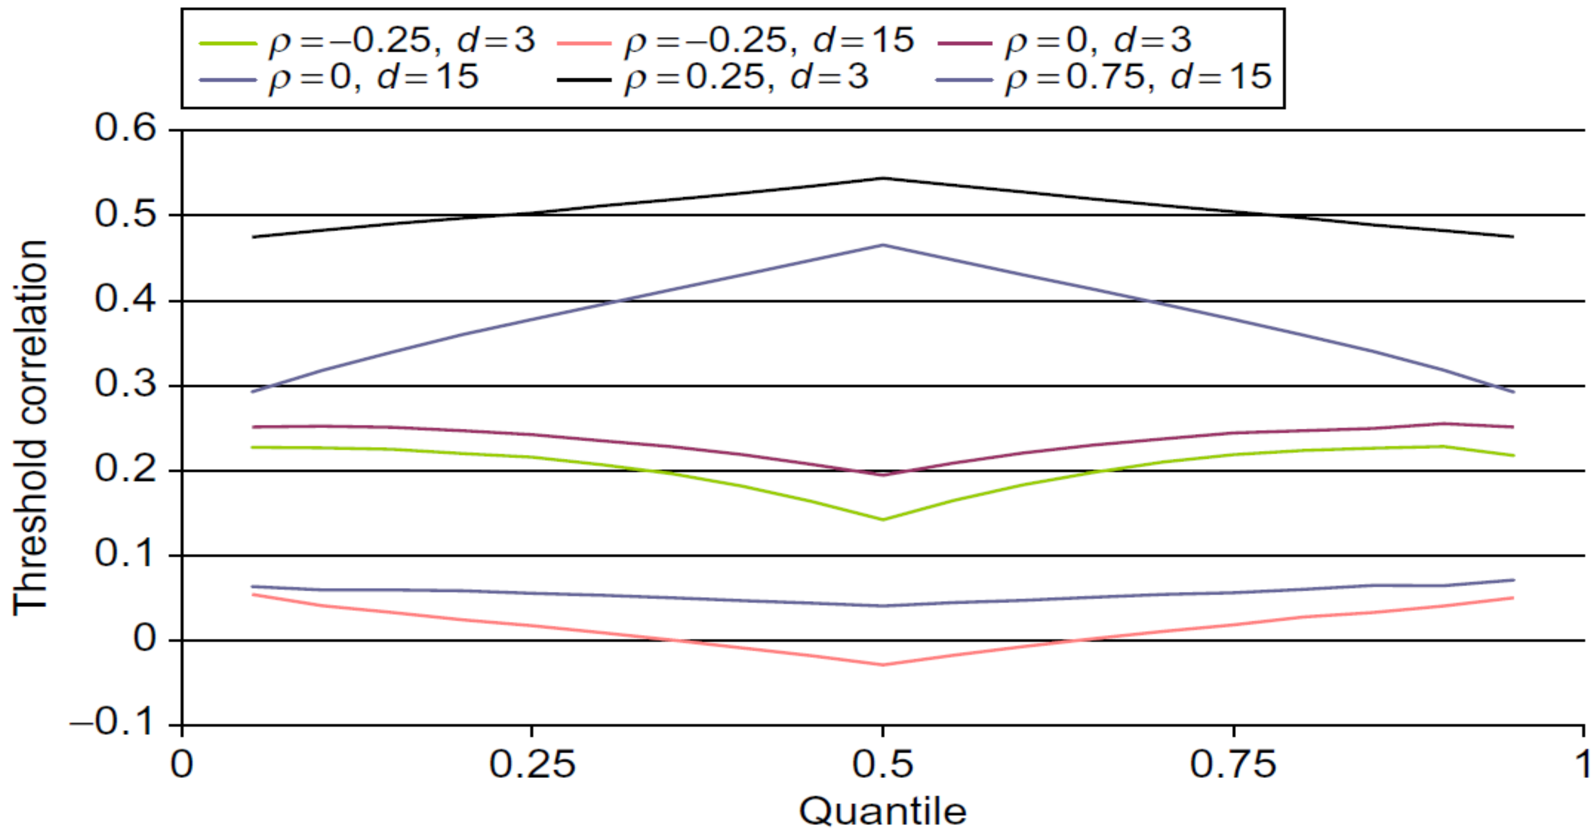

$$G(u_1,\ldots,u_n;\Upsilon^*,d) = t_{(d,\Upsilon^*)}\left(t^{-1}(u_1;d),\ldots,t^{-1}(u_n;d)\right)$$

$$g(u_1,\ldots,u_n;\Upsilon^*,d) = \frac{t_{(d,\Upsilon^*)}\left(t^{-1}(u_1;d),\ldots,t^{-1}(u_n;d)\right)}{\prod_{i=1}^n t(t^{-1}(u_i;d);d)}$$

$$= \frac{\Gamma\left(\frac{d+n}{2}\right)}{|\Upsilon^*|^{\frac{1}{2}}\Gamma\left(\frac{d}{2}\right)}\left(\frac{\Gamma\left(\frac{d}{2}\right)}{\Gamma\left(\frac{d+1}{2}\right)}\right)^n \frac{\left(1+\frac{1}{d}t^{-1}(u;d)'\Upsilon^{*-1}t^{-1}(u;d)\right)^{-\frac{d+n}{2}}}{\prod_{i=1}^n\left(1+\frac{(t^{-1}(u_i;d))^2}{d}\right)^{-\frac{d+1}{2}}}$$

- An asymmetric t copula can be developed from the asymmetric multivariate t distribution in the same way

(Threshold) Correlations with Copulas : Pros and Cons

⊕ The full flexibility allowed by Sklar's theorem that allows us to "build" multivariate models by composition of univariate models

⊕ Specific attention is devoted to picking copulas that fit the correlations in the tails, which is related to the notions underlying stress scenarios

⊖ Fails to nest any other key econometric framework

⊖ Unclear how to pick copulas among many alternatives

⊖ No economic intuition for the choice of copulas

⊖ Unclear (or awkward) how to condition copulas on time-varying information flows:

$$F_{\mathcal{F}_t}(z_1, \ldots, z_n) = G_{\mathcal{F}_t}(F_1(z_1), \ldots, F_n(z_n)) = G_{\mathcal{F}_t}(u_1, \ldots, u_n)$$