Generalised Covariances and Correlations

Research Seminar
Institute for Statistics and Mathematics
Vienna University of Economics and Business

Tobias Fissler[1]
Marc-Oliver Pohle[2]

2nd December 2022

1 Vienna University of Economics and Business
2 Goethe University Frankfurt

Introduction

Dependence

- Measuring and modelling dependence between random variables is at the heart of statistics and of paramount importance in every empirical discipline.
- Two key approaches:
    - (Mutual) dependence measures, e.g. Pearson correlation
    - Regression, e.g. least-squares regression

Measures of Dependence

- We are interested in the mutual dependence between two random variables X and Y (direction and strength of dependence).
- Joint CDF: $F_{X,Y}$
- Marginal CDFs: F_X, F_Y; assume continuity for this talk
- A measure of dependence for X and Y, $\delta(X, Y) := \delta(F_{X,Y})$, maps $F_{X,Y}$ to a real vector space, usually the real line.

Desirable Properties

Going back to Rényi's axioms (Rényi, 1959) for measures of dependence, often modified (see e.g. Schweizer and Wolff (1981); Embrechts et al. (2002); Balakrishnan and Lai (2009)):

- Independence: $\delta(X, Y) = 0$ if X and Y are independent.
- Normalisation: $\delta(X, Y) \in [-1, 1]$.
- Attainability: $\delta(X, Y) = \pm 1$ if X and Y have perfect positive (negative) dependence, i.e. are comonotonic (countermonotonic): $Y = g(X)$ for some increasing (decreasing) g.
- Invariance to strictly increasing transformations: Let g be a strictly increasing function: $\delta(g(X), Y) = \delta(X, g(Y)) = \delta(X, Y)$ ($\iff$ invariance to marginals F_X, F_Y)
- Sometimes: symmetry ($\delta(X, Y) = \delta(Y, X)$) and the reverse direction for independence and attainability

Literature on Measures of Dependence

- Pearson correlation is the most widely-used measure of dependence, but has major shortcomings, in particular non-attainability (see e.g. Embrechts et al. (2002)).
- Rank correlations, i.e. Spearman's ρ and Kendall's τ are popular alternatives with nice properties.
- Many other measures have been proposed (see e.g. Balakrishnan and Lai (2009); Tjøstheim et al. (2022) for overviews).

Motivation i: Beyond Average Dependence

- Usually those are scalar-valued measures of global dependence, trying to summarise overall/average dependence between X and Y in a single number $\rightarrow$ very limited information on dependence structure contained in $F_{X,Y}$
- Measures of local dependence?
- Measures that characterize full dependence structure?
- Work going in that direction: Holland and Wang (1987), Tjøstheim and Hufthammer (2013)

Example: Center vs. Tails

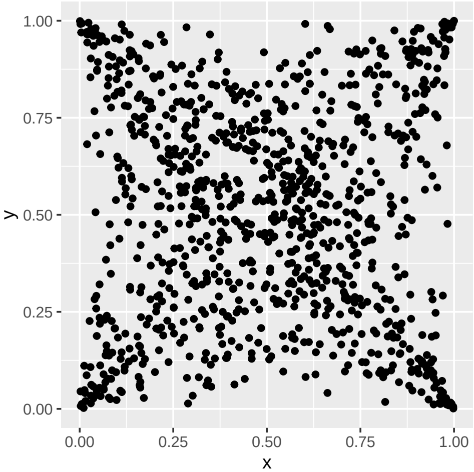

Figure 1: scatter plot of 1000 draws from a bivariate Cauchy copula with Spearman's ρ equal to 0

Recap: Pearson Correlation

Covariance and Pearson correlation

$$\mathrm{Cov}(X, Y) = \mathrm{E}[(X - \mu(X))(Y - \mu(Y))], \; r(X, Y) = \frac{\mathrm{Cov}(X, Y)}{\sqrt{\mathrm{Var}(X)\,\mathrm{Var}(Y)}}$$

They measure co-movements of X and Y around their means.

Motivation ii: Quantile or even Generalised Correlation?

- Can we also measure co-movements around quantiles or other statistical functionals?
- Quantile or generalised regression generalises least-squares regression. Generalisation of Pearson correlation along those lines?
- Starting point for this project
- Some proposals for quantile correlations exist, e.g. Linton and Whang (2007); Li et al. (2015)

This Paper i

Generalised correlation:

- Measures co-movements of X and Y around functionals other than the mean (e.g. quantiles, expectiles).
- Two key ingredients: generalised covariance and normalisation
- Desirable theoretical properties
- Many useful measures arise as special cases
- In particular local correlations: quantile and threshold correlation, allowing to measure local dependence
- Also mean correlation, an improved version of Pearson correlation

This Paper ii

Distributional covariances and correlations:

- Two closely related measures: quantile function and CDF correlation (CDF: cumulative distribution function)
- Basically generalised correlation with quantile function and CDF as functionals
- Amounts to considering the set of all quantile or threshold correlations jointly
- Not numbers, but functions on $[0,1]^2$ and $\mathbb{R}^2$
- Uncover the full dependence structure between X and Y
- Desirable properties
- Close relationship to copula and joint CDF $F_{X,Y}$
- Lead to clear definitions of positive and negative dependence

This Paper iii

Tail correlation:

- New measure of tail dependence arising as a limit of quantile correlation.
- Closely related and improving upon well-known coefficient of tail dependence.

Summary covariances and correlations:

- If the goal is to summarize dependence in one number: Natural to integrate over distributional covariances w.r.t. an arbitrary measure and then normalize
- Spearman's ρ and mean correlation arise as special cases.
- Other useful measures can be constructed.

Generalised Covariance

Statistical Functionals

- A statistical functional is a map $T : \mathcal{M} \to \mathbb{R}$ where $\mathcal{M}$ is a class of probability distributions.
- Examples: Mean, median, expectiles, quantiles, exceedance probabilities, ...
- We can evaluate T for random variables by setting $T(X) := T(F_X)$ for $X \sim F_X$.

Generalised Covariance?

- Recall classical covariance:
  $\mathrm{Cov}(X, Y) = \mathrm{E}[(X - \mu(X))(Y - \mu(Y))]$
- Independence implies nullity due to the fact that
  $\mathrm{E}[X - \mu(X)] = \mathrm{E}[Y - \mu(Y)] = 0$.
- Just replacing μ by some other functional T violates this.
- We need a suitable way to express deviations of a random variable X from a functional $T(X)$, that is, a generalised error to replace $e_\mu(X) = X - \mu(X)$.

We employ identification functions to define generalised errors.

Definition (Identification function)

A map $v : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is an $\mathcal{M}$-identification function for T if $\int |v(t,x)| \mathrm{d}F(x) < \infty$ for all $t \in \mathbb{R}$, $F \in \mathcal{M}$ and if

$$t = T(F) \quad \Longrightarrow \quad \int v(t,x) \mathrm{d}F(x) = 0$$

for all $t \in \mathbb{R}$, $F \in \mathcal{M}$. It is a strict $\mathcal{M}$-identification function if $\Longleftarrow$ holds as well.

Identification Functions ii

Examples:

- Mean $T = \mu$: $v(t, x) = x - t$
- α-quantile $T = q_\alpha$: $v(t, x) = \alpha - \mathbb{1}\{x \leq t\}$
- Exceedance probability $T(F) = F(a)$: $v(t, x) = t - \mathbb{1}\{x \leq a\}$.

Usual Uses:

- In forecast evaluation, identification functions are used to check calibration of forecasts.
- In econometrics, they are often called moment functions.
- In estimation, they give rise to Z-estimation.

Generalised Error

Definition (Generalised error)

Let $v(t, x)$ be an identification function for $T : \mathcal{M} \to \mathbb{R}$ which is increasing in x for all t. Then

$$e_T(X) = v(T(X), X)$$

is the generalised error of X for T.

Examples:

- Mean $T = \mu$: $e_\mu(X) = X - \mu(X)$
- α-quantile $T = q_\alpha$: $e_{q_\alpha}(X) = \alpha - \mathbb{1}\{X \leq q_\alpha(X)\}$
- Exceedance probability $T(F) = F(a)$:
  $e_{F(a)}(X) = F(a) - \mathbb{1}\{X \leq a\}$.

Properties of the Generalised Error

$e_T(X)$ is a suitable way to express deviations of X from $T(X)$:

- It is centred: $\mathrm{E}[e_T(X)] = 0$.
- It has the right sign.
- It gets larger when X is further away from T.

Generalised Covariance

Definition (Generalised covariance)

Let e_{T_1}, e_{T_2} be generalised errors for T_1 and T_2. Then the generalised covariance at T_1 and T_2 is

$$\mathrm{Cov}_{T_1,T_2}(X, Y) = \mathrm{E}\left[e_{T_1}(X)e_{T_2}(Y)\right].$$

- Cov_{T_1,T_2} measures average co-movements of X and Y around their respective functionals.
- It holds that $\mathrm{Cov}_{T_1,T_2}(X, Y) = \mathrm{Cov}\left(e_{T_1}(X), e_{T_2}(Y)\right)$.
- $\mathrm{Cov}_{T_1,T_2}(X, Y) = 0$ if X and Y are independent.
- Cov_{T_1,T_2} depends on the choices of the identification functions, we use the canonical ones from Gneiting and Resin (2021). This dependence disappears for the generalised correlation.

Generalised Correlation

Normalisation: Cauchy-Schwarz

- The normalisation in Pearson correlation exploits the Cauchy–Schwarz inequality:

$$\left| \mathrm{E}\left[e_{T_1}(X) e_{T_2}(Y)\right] \right| \leq \sqrt{\mathrm{E}[e_{T_1}(X)^2]\ \mathrm{E}[e_{T_2}(Y)^2]}.$$

- Hence, the Cauchy–Schwarz normalisation ensures that

$$\frac{\mathrm{Cov}_{T_1,T_2}(X,Y)}{\sqrt{\mathrm{E}[e_{T_1}(X)^2]\ \mathrm{E}[e_{T_2}(Y)^2]}} \in [-1,1].$$

- Problem with attainability (inherited from Pearson correlation): This quantity may have a maximum (minimum) that is far away from 1 (-1) and their absolute values may be very different.

Normalisation: Co- and Countermonotonicity

- Solution: We need a sharp normalisation that distinguishes between positive and negative dependence.
- Alternative inequality:

$$\mathrm{E}\left[e_{T_1}(X)e_{T_2}(Y')\right] \leq \mathrm{E}\left[e_{T_1}(X)e_{T_2}(Y)\right] \leq \mathrm{E}\left[e_{T_1}(X)e_{T_2}(Y'')\right]$$

  where Y, Y' and $Y'' \sim F_Y$ and where
  - (X, Y') are countermonotonic;
  - (X, Y'') are comonotonic.
- This inequality is sharp by construction.

Generalised Correlation

Definition (Generalised Correlation)

Let Cov_{T_1,T_2} be a generalised covariance at T_1 and T_2. Then the corresponding generalised correlation is

$$\mathrm{Cor}_{T_1,T_2}(X,Y) = \begin{cases} \frac{\mathrm{Cov}_{T_1,T_2}(X,Y)}{|\mathrm{Cov}_{T_1,T_2}(X,Y')|}, & \text{if } \mathrm{Cov}_{T_1,T_2}(X,Y) \leq 0 \\ \frac{\mathrm{Cov}_{T_1,T_2}(X,Y)}{|\mathrm{Cov}_{T_1,T_2}(X,Y'')|}, & \text{if } \mathrm{Cov}_{T_1,T_2}(X,Y) > 0 \end{cases}$$

- $\mathrm{Cor}_{T_1,T_2}(X,Y)$ fulfills independence, normalisation and attainability.
- Symmetry if $T_1 = T_2$.
- Invariance to increasing transformations depends on choice of functionals T_1 and T_2.

Examples: Quantile Correlation i

- Quantiles $T_1 = q_\alpha$, $T_2 = q_\beta$:

$$\mathrm{QCov}_{\alpha,\beta}(X, Y) = E\big[(\alpha - \mathbb{1}\{X \leq q_{\alpha(X)}\})(\beta - \mathbb{1}\{X \leq q_{\beta(X)}\})\big],$$

- Normalisation:

$$\mathrm{QCov}_{\alpha,\beta}(X, Y') = \max(\alpha + \beta - 1, 0) - \alpha\beta,$$
$$\mathrm{QCov}_{\alpha,\beta}(X, Y'') = \min(\alpha, \beta) - \alpha\beta.$$

(Fréchet-Hoeffding bounds)

- Relation to the copula of (X, Y), $C_{X,Y}$:

$$\mathrm{QCov}_{\alpha,\beta}(X, Y) = C_{X,Y}(\alpha, \beta) - \alpha\beta$$

- Does not depend on the marginals F_X and F_Y, i.e. is invariant to strictly increasing transformations

Examples: Quantile Correlation ii

- Measures local dependence, e.g. tail dependence
- Connection to quantilogram (Linton and Whang, 2007): Quantile Covariance (with same α) with Cauchy–Schwarz normalisation. Or the cross-quantilogram of Han et al. (2016).
- Special case: median correlation with $\alpha = \beta = 0.5$, equivalent to Blomqvist's β (Blomqvist, 1950)

Examples: Threshold Correlation

- Threshold covariance, $T_1(F) = F(a)$, $T_2(F) = F(b)$:

$$\mathrm{TCov}_{a,b}(X, Y) = E\big[(F(a) - \mathbb{1}\{X \leq a\})(F(b) - \mathbb{1}\{Y \leq b\})\big]$$

- Normalisation:

$$\mathrm{TCov}_{a,b}(X, Y') = \max(F_X(a) + F_Y(b) - 1, 0) - F_X(a)F_Y(b),$$
$$\mathrm{TCov}_{a,b}(X, Y'') = \min(F_X(a), F_Y(b)) - F_X(a)F_Y(b).$$

  (Fréchet-Hoeffding bounds)

- Relation to joint and marginal CDFs:

$$\mathrm{TCov}_{a,b}(X, Y) = F_{X,Y}(a, b) - F_X(a)F_Y(b).$$

- Allows to measure local dependence as well

Local Correlations

- Quantile and threshold correlation
- Natural complements: measuring dependence on the quantile scale (and independent of marginals) or on the observation scale
- Building blocks for further measures: distributional, tail and summary correlations

Examples: Mean Correlation

- Mean $T_1 = T_2 = \mu$: $\mathrm{Cov}_\mu = \mathrm{Cov}(X, Y)$. But Cor_μ is generally different from Pearson correlation Cor due to the normalisation.
- Normalisation: $Cov(X, Y')$ and $Cov(X, Y'')$ instead of $\sqrt{\mathrm{Var}(X)\,\mathrm{Var}(Y)}$
- Mean correlation can be seen as an improved version of Pearson correlation, ensuring attainability.
- In cases, where Pearson correlation is attainable (e.g. under multivariate normality), the two coincide.

Examples: Mean-quantile Correlation

- Quantile and mean, $T_1 = q_\alpha$, $T_2 = \mu$:

$$\mathrm{Cov}_{q_\alpha,\mu}(X, Y) = \mathrm{E}\left[(\alpha - \mathbb{1}\{X \leq q_\alpha(X)\})(Y - \mu(Y))\right]$$

- Connection to what Li et al. (2015) call quantile correlation, where Cauchy-Schwarz normalisation is used.

Estimation

Suppose an iid sample (X_i, Y_i), $i = 1, \dots, n$, is available.

- Estimate T_1 and T_2 by empirical counterparts $\widehat{T}_1(X) =: \hat{t}_1$ and $\widehat{T}_2(Y) =: \hat{t}_2$ and $\mathrm{Cov}_{T_1,T_2}(X, Y)$ by

$$\widehat{\mathrm{Cov}}_{T_1,T_2}(X, Y) = \frac{1}{n}\sum_{i=1}^{n} v_{T_1}(\hat{t}_1, X_i) v_{T_2}(\hat{t}_2, Y_i).$$

- For Hoeffding normalisation, use order statistics: $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$ and $Y_{(1)} \leq Y_{(2)} \leq \cdots \leq Y_{(n)}$.

$$\widehat{\mathrm{Cov}}_{T_1,T_2}(X, Y'') = \frac{1}{n}\sum_{i=1}^{n} v_{T_1}(\hat{t}_1, X_{(i)}) v_{T_2}(\hat{t}_2, Y_{(i)}),$$

$$\widehat{\mathrm{Cov}}_{T_1,T_2}(X, Y') = \frac{1}{n}\sum_{i=1}^{n} v_{T_1}(\hat{t}_1, X_{(i)}) v_{T_2}(\hat{t}_2, Y_{(n-i+1)}).$$

Asymptotics

- $\widehat{\mathrm{Cov}}_{T_1,T_2}$ and $\widehat{\mathrm{Cor}}_{T_1,T_2}$ are strongly consistent (under mild conditions).
- Conjecture on the limiting distribution:
  - For $\mathrm{Cov}_{T_1,T_2} \neq 0$: normal
  - For $\mathrm{Cov}_{T_1,T_2} = 0$: combination of two halves of normals with mean zero and different variances

Distributional Covariances and Correlations

Distributional Correlations: Definition

- Idea: Look at all the local information jointly, i.e. consider quantile correlation at all combinations of quantile levels $(\alpha, \beta) \in (0,1)^2$ or threshold correlation at all $(x, y) \in \mathbb{R}^2$.
- Define distributional correlations accordingly,CDF correlation and quantile function correlation:

$$\mathrm{CDFCor}(X, Y) = \big(\,\mathrm{TCor}_{a,b}(X, Y)\big)_{a,b\in\mathbb{R}}$$

$$\mathrm{QFCor}(X, Y) = \big(\,\mathrm{QCor}_{\alpha,\beta}(X, Y)\big)_{\alpha,\beta\in[0,1]}$$

- Those dependence measures are not numbers, but two-dimensional functions.
- There is an alternative way to arrive at them...

Distributional Correlations: Alternative Definition

- One could define generalised covariances for vector- and function-valued functionals via the outer product of vector- and function-valued generalised errors induced by corresponding identification functions.
- Choosing the CDF and the quantile function themselves as functionals,

$$T_{CDF}(F) = F, \qquad T_{QF}(F) = F^{-1},$$

  with the identification functions

$$v_{CDF}(F, x) = \big(F(a) - \mathbb{1}\{x \leq a\}\big)_{a \in \mathbb{R}},$$

$$v_{QF}(F^{-1}, x) = \big(\alpha - \mathbb{1}\{x \leq F^{-1}(\alpha)\}\big)_{\alpha \in [0,1]}$$

  would also lead to the two distributional covariances.
- Normalise pointwise to arrive at the respective correlations.

Distributional Correlations: Relation to Copula and Joint CDF i

- Distributional correlations indeed uncover the full dependence structure between X and Y as the following representations in terms of copula and joint CDF show.
- For quantile correlation:

$$\mathrm{QFCor}(X, Y; \alpha, \beta) = \begin{cases} \frac{C_{X,Y}(\alpha,\beta)-\alpha\beta}{\min(\alpha,\beta)-\alpha\beta}, & \mathrm{QFCov}(X, Y; \alpha, \beta) \geq 0 \\ \frac{C_{X,Y}(\alpha,\beta)-\alpha\beta}{-\max(\alpha+\beta-1,0)+\alpha\beta}, & \mathrm{QFCov}(X, Y; \alpha, \beta) < 0 \end{cases}.$$

- Note that here the three limiting cases from copula theory, namely the independence, the co- and the countermonotonicity copula, show up.

Distributional Correlations: Relation to Copula and Joint CDF ii

- For CDF correlation, we have

$$\begin{aligned} &\mathrm{CDFCor}(X, Y; a, b) \\ &= \begin{cases} \frac{F_{X,Y}(a,b) - F_X(a)F_Y(b)}{\min(F_X(a), F_Y(b)) - F_X(a)F_Y(b)}, & \mathrm{CDFCov}(X, Y; a, b) \geq 0 \\ \frac{F_{X,Y}(a,b) - F_X(a)F_Y(b)}{-\max(F_X(a) + F_Y(b) - 1, 0) + F_X(a)F_Y(b)}, & \mathrm{CDFCov}(X, Y; a, b) < 0 \end{cases}. \end{aligned}$$

- Thus, distributional correlations are closely related to copula and joint CDF and make the dependence structure contained in them explicit and visible.

Distributional Correlations: Properties

They inherit the properties from the respective local correlations. Further:

Proposition

(i) *X and Y are independent if and only if* $\mathrm{QFCor}(X, Y) = 0$.

(ii) *X and Y are comonotonic if and only if* $\mathrm{QFCor}(X, Y) = 1$.

(iii) *X and Y are countermonotonic if and only if* $\mathrm{QFCor}(X, Y) = -1$.

The same results hold for CDFCor.

Example: QFCor for a bivariate Cauchy copula

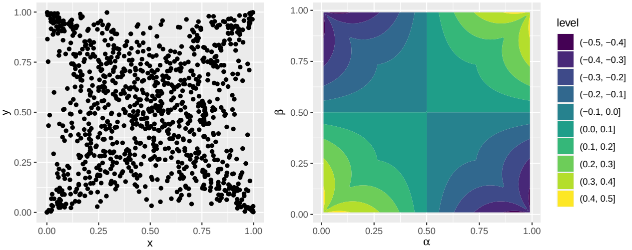

Figure 2: plot of QFCor for a bivariate Cauchy copula with Spearman's ρ equal to 0 and scatter plot of 1000 draws from it

Examples: Normal, Cauchy, Clayton, Gumbel copula

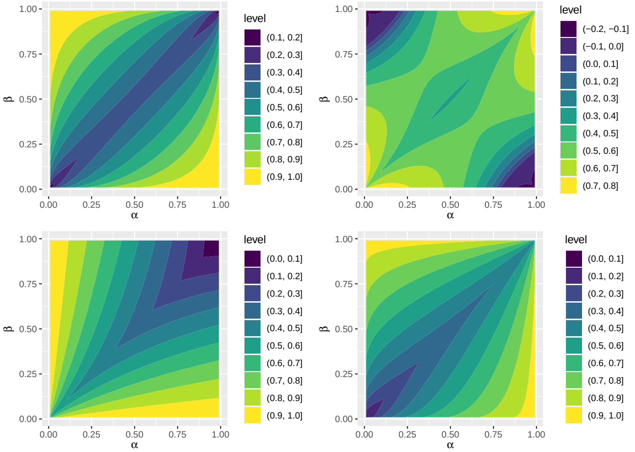

Figure 3: plots of QFCor for four different copulas, all with Spearman's ρ equal to 0.5: Gaussian (upper left), Cauchy (upper right), Clayton (lower left) and Gumbel (lower right)

Defining Positive and Negative Dependence

- How to define positive and negative dependence between two random variables X and Y?
- Countless proposals in the literature, see e.g. Balakrishnan and Lai (2009) for an overview.
- Distributional correlations suggest a natural definition:

Definition

X and Y are positively (negatively) dependent if

$$\mathrm{CDFCor}(X, Y) \geq 0 \qquad (\mathrm{CDFCor}(X, Y) \leq 0).$$

- Positive (negative) dependence correspond to positive (negative) quadrant dependence due to Lehmann (1966).
- Using *QFCor* leads to an equivalent definition.

A Further Property of Generalised Covariance and Correlation

Proposition

If X and Y are positively (negatively) dependent, then for any T_1, T_2

$$\mathrm{Cov}_{T_1,T_2}(X,Y) \geq 0 \qquad (\mathrm{Cov}_{T_1,T_2}(X,Y) \leq 0).$$

Tail Correlation

Tail Correlation

- It is often of interest to analyse co-movement in the tails.
- Quantile correlation suggests a natural measure of tail dependence.

Definition (Tail correlations)

The lower tail correlation is defined as

$$\mathrm{LTCor}(X, Y) := \lim_{\alpha \to 0} \mathrm{QCor}_{\alpha,\alpha}(X, Y).$$

The upper tail correlation is defined as

$$\mathrm{UTCor}(X, Y) := \lim_{\alpha \to 1} \mathrm{QCor}_{\alpha,\alpha}(X, Y).$$

Relation to Coefficient of Tail Dependence

- The by far most prominent tail dependence measure is the tail dependence coefficient (see e.g. Joe (1993) or Coles et al. (1999)).
- Consider e.g. the lower tail dependence coefficient:

$$\lambda_l(X,Y) := \lim_{\alpha\to 0} \mathrm{P}(Y \leq q_\alpha(Y) | X \leq q_\alpha(X)).$$

- Under positive dependence, i.e. if there is some $\alpha_0 \in (0,1)$ such that $\mathrm{QCov}_{\alpha,\alpha}(X,Y) \geq 0$ for all $\alpha \in (0,\alpha_0)$, we have

$$\mathrm{LTCor}(X,Y) = \lambda_l(X,Y).$$

- However, $\lambda_l(X,Y) = 0$ under negative dependence and independence, while $\mathrm{LTCor}(X,Y)$ shows the desirable behaviour, i.e. is 0 under independence, negative under negative dependence and -1 under countermonotonicity.

Summary Covariances and Correlations

Summary Correlations: Idea

- Idea: Summarize the full dependence structure contained in the distributional correlations via a single number by integration.
- Take $\mathrm{QFCov}(X, Y) = \big(\mathrm{QCov}_{\alpha,\beta}(X, Y)\big)_{\alpha,\beta\in[0,1]}$ and integrate it with respect to a measure π on $[0, 1]^2$:

$$\int_{[0,1]^2} \mathrm{QCov}_{\alpha,\beta}(X, Y)\,\mathrm{d}\pi(\alpha, \beta)$$

- Likewise, take $\mathrm{CDFCov}(X, Y)$ and integrate it with respect to a measure on $\mathbb{R}^2$.
- Depending on the sign of those integrals, normalise with the integrals over the corresponding Fréchet-Hoeffding bound.

Summary correlations inherit properties of respective local correlations. Further:

Proposition

The summary correlations fulfil:

- *They are 1 if and only if X and Y are comonotonic.*
- *They are -1 if and only if X and Y are countermonotonic.*

Proposition

(i) *If* $\mathrm{CDFCov}(X, Y)$ *is integrated with respect to the Lebesgue measure on* $\mathbb{R}^2$*, we obtain covariance. Hence, the respective summary correlation retrieves mean correlation (or Pearson correlation if the Cauchy-Schwarz normalisation is used).*

(ii) *If* $\mathrm{QFCov}(X, Y)$ *is integrated with respect to the Lebesgue measure on* $[0, 1]^2$ *and normalised, we obtain Spearman's* ρ*.*

This justifies the use of Spearman's ρ and mean/ Pearson correlation as measures of overall/average dependence.

New Measures of Tail Dependence

By integrating only over certain subsets, one can obtain further interesting summary correlations, e.g. new measures of tail dependence.

Conclusion

Summary i: Generalised Correlation

- Generalised correlation: Measures average co-movements around general functionals.
- Similar paradigm shift as from mean regression to quantile or generalised regression.
- Nice theoretical properties
- Several interesting measures arise

Summary ii: A Measure for Every Purpose

- Two families of correlations containing a measure for every purpose: a local, a distributional and a summary correlation.
- The quantile family: quantile correlation, quantile function correlation and Spearman's ρ.
- The CDF family: threshold, CDF and mean correlation.
- Very natural, closely related to fundamental statistical concepts, i.e. copula and CDF.

Outlook

- This paper: applications
- Future work: inference

References

Narayanaswamy Balakrishnan and Chin Diew Lai. *Continuous bivariate distributions*. Springer Science & Business Media, 2009.

Nils Blomqvist. On a measure of dependence between two random variables. *The Annals of Mathematical Statistics*, pages 593–600, 1950.

Stuart Coles, Janet Heffernan, and Jonathan Tawn. Dependence measures for extreme value analyses. *Extremes*, 2(4):339–365, 1999.

Paul Embrechts, Alexander McNeil, and Daniel Straumann. Correlation and dependence in risk management: properties and pitfalls. In M. Dempster, editor, *Risk management: value at risk and beyond*, pages 176–223. Cambridge University Press, 2002.

Tilmann Gneiting and Johannes Resin. Regression diagnostics meets forecast evaluation: Conditional calibration, reliability diagrams, and coefficient of determination. *arXiv preprint arXiv:2108.03210*, 2021.

Paul W Holland and Yuchung J Wang. Dependence function for continuous bivariate densities. *Communications in Statistics-Theory and Methods*, 16(3):863–876, 1987.

Harry Joe. Parametric families of multivariate distributions with given margins. *Journal of multivariate analysis*, 46(2): 262–282, 1993.

Erich Leo Lehmann. Some concepts of dependence. *The Annals of Mathematical Statistics*, 37(5):1137–1153, 1966.

Guodong Li, Yang Li, and Chih-Ling Tsai. Quantile correlations and quantile autoregressive modeling. *Journal of the American Statistical Association*, 110(509):246–261, 2015.

Oliver Linton and Yoon-Jae Whang. The quantilogram: With an application to evaluating directional predictability. *Journal of Econometrics*, 141(1):250–282, 2007.

Alfréd Rényi. On measures of dependence. *Acta mathematica hungarica*, 10(3-4):441–451, 1959.

Berthold Schweizer and Edward F Wolff. On nonparametric measures of dependence for random variables. *The annals of statistics*, 9(4):879–885, 1981.

Dag Tjøstheim and Karl Ove Hufthammer. Local Gaussian correlation: A new measure of dependence. *Journal of Econometrics*, 172(1):33–48, 2013.

Dag Tjøstheim, Håkon Otneim, and Bård Støve. Statistical dependence: Beyond Pearson's rho. *Statistical Science*, 37(1): 90 – 109, 2022. doi: 10.1214/21-STS823. URL *https://doi.org/10.1214/21-STS823*.

Recap: Copulas

- bivariate CDF of RVs X and Y: $F_{X,Y}(x,y) = \mathrm{P}(X \leq x, Y \leq y)$
- copula: CDF with $U(0,1)$ margins
- Sklar's theorem:
  - $F_{X,Y}$ can be written as $F_{X,Y}(x,y) = C_{X,Y}(F_X(x), F_Y(y))$, where F_X and F_Y are the marginal distributions.
  - $C_{X,Y}(u,v) = F(F_X^{-1}(u), F_Y^{-1}(v))$

Attainability Lemma

Lemma (Attainability of Pearson correlation)

Let $X, Y \in L^2(\mathbb{R})$ be non-constant. Pearson correlation $\mathrm{Cor}(X, Y)$ is attainable, that is, there exist joint distributions $F, \tilde{F}$ with marginals F_X, F_Y and Pearson correlation 1 and -1, respectively, if and only if X and Y are of the same type, that is, $F_Y = F_{a+bX}$ for some $b > 0$ and $a \in \mathbb{R}$, and the distributions are symmetric, that is, there exist $c, d \in \mathbb{R}$ such that $F_{c+X} = F_{c-X}$ and $F_{d+Y} = F_{d-Y}$.

Attainability Problem of the CS Normalization for QCov

For example, it can be shown that

$$r\big(e_{q_\alpha}(X), e_{q_\alpha}(Y)\big) \geq \big(\max(2\alpha - 1, 0) - \alpha^2\big)/\big(\alpha(1-\alpha)\big).$$

While this yields a desirable lower bound of -1 for $\alpha = 1/2$, the lower bound converges to 0 as α gets closer to 0 or 1.

E.g., for $\alpha = 0.95$ the lower bound is approximately -0.05, for $\alpha = 0.75$ it is $-1/3$.