Copula theory and applications: Part I

Hans Colonius

Department of Psychology
Oldenburg University
hans.colonius@uni-oldenburg.de
https://uol.de/en/hans-colonius

Outline

Example 1

(i) If (Z_1, Z_2) follows a bivariate Gaussian distribution, do Z_1 and Z_2 follow univariate Gaussian distributions?

Example 1

(i) If (Z_1, Z_2) follows a bivariate Gaussian distribution, do Z_1 and Z_2 follow univariate Gaussian distributions?

(ii) If Z_1 and Z_2 follow univariate Gaussian distributions, does (Z_1, Z_2) follow a bivariate Gaussian distribution?

Example 1

(i) If (Z_1, Z_2) follows a bivariate Gaussian distribution, do Z_1 and Z_2 follow univariate Gaussian distributions?

(ii) If Z_1 and Z_2 follow univariate Gaussian distributions, does (Z_1, Z_2) follow a bivariate Gaussian distribution?

COPULAS ALLOW TO MODEL MULTIVARIATE (STOCHASTIC) DEPENDENCY SEPARATELY FROM THE MARGINALS

Some important definitions

1. F is non-decreasing: $x_1 < x_2 \implies F(x_1) \leq F(x_2)$;
2. F is right-continuous: $\lim_{h \to 0^+} F(x+h) = F(x)$;
3. $\lim_{x \to -\infty} F(x) = 0$ and $\lim_{x \to \infty} F(x) = 1$;
4. F has at most a countable number of discontinuities.

- If a function $F : \mathbb{R} \to [0, 1]$ is non-decreasing, right-continuous, $\lim_{x \to -\infty} F(x) = 0$, and $\lim_{x \to \infty} F(x) = 1$, then F is the **distribution function** of some random variable.

Distribution function

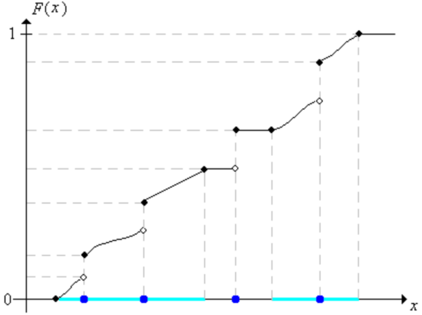

Some important definitions

A distribution function F is called

- **discrete** if for some countable set of numbers $\{x_j\}$ and point masses $\{p_j\}$,

$$F(x) = \sum_{x_j \leq x} p_j, \quad \text{for all } x \in \mathbb{R}.$$

The function p is called **probability mass function**;

Some important definitions

A distribution function F is called

- **discrete** if for some countable set of numbers $\{x_j\}$ and point masses $\{p_j\}$,

$$F(x) = \sum_{x_j \leq x} p_j, \quad \text{for all } x \in \mathbb{R}.$$

The function p is called **probability mass function**;

- **absolutely continuous** if there exists a non-negative, (Lebesgue) integrable function f, such that

$$F(b) - F(a) = \int_a^b f(x)\,dx \quad \text{for all } a < b.$$

The function f is called **density** of F;

Some important definitions

A <u>distribution function</u> F is called

- **discrete** if for some countable set of numbers $\{x_j\}$ and point masses $\{p_j\}$,

$$F(x) = \sum_{x_j \leq x} p_j, \text{ for all } x \in \mathbb{R}.$$

  The function p is called **probability mass function**;

- **absolutely continuous** if there exists a non-negative, (Lebesgue) integrable function f, such that

$$F(b) - F(a) = \int_a^b f(x)\,dx \text{ for all } a < b.$$

  The function f is called **density** of F;

- **singular** if $F \neq 0, F'$ exists and equals 0 a.e. (almost everywhere).

Discrete distribution function

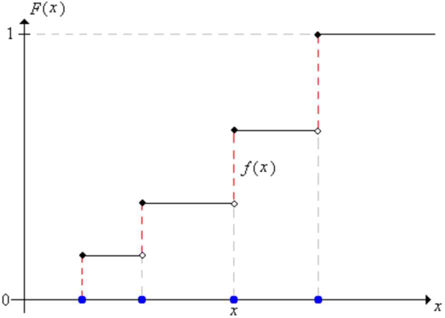

Continuous distribution function

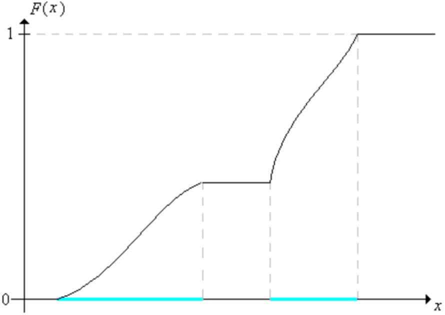

Some important definitions

The **joint, or multivariate, distribution function** of random vector $\mathbf{X} = (X_1, X_2, \ldots, X_n)'$

$$F_{X_1,X_2,\ldots,X_n}(x_1,, x_2, \ldots, x_n) = P(X_1 \leq x_1, X_2 \leq x_2, \ldots, X_n \leq x_n),$$

for $x_k \in \mathbb{R}$, $k = 1, 2, \ldots, n$. This is written more compactly as

$$F_{\mathbf{X}}(\mathbf{x}) = P(\mathbf{X} \leq \mathbf{x})$$

.

Some important definitions

The **joint, or multivariate, distribution function** of random vector $\mathbf{X} = (X_1, X_2, \ldots, X_n)'$

$$F_{X_1,X_2,\ldots,X_n}(x_1,, x_2, \ldots, x_n) = P(X_1 \leq x_1, X_2 \leq x_2, \ldots, X_n \leq x_n),$$

for $x_k \in \mathbb{R}$, $k = 1, 2, \ldots, n$. This is written more compactly as

$$F_{\mathbf{X}}(\mathbf{x}) = P(\mathbf{X} \leq \mathbf{x})$$

.

- For discrete distributions the **joint probability mass function** is defined by

$$p_{\mathbf{X}}(\mathbf{x}) = P(\mathbf{X} = \mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^n$$

- and in the absolutely continuous case we have a **joint density**

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{\partial^n F_{\mathbf{X}}(\mathbf{x})}{\partial x_1 \partial x_2 \cdots \partial x_n}, \quad \mathbf{x} \in \mathbb{R}^n.$$

Bivariate normal density function

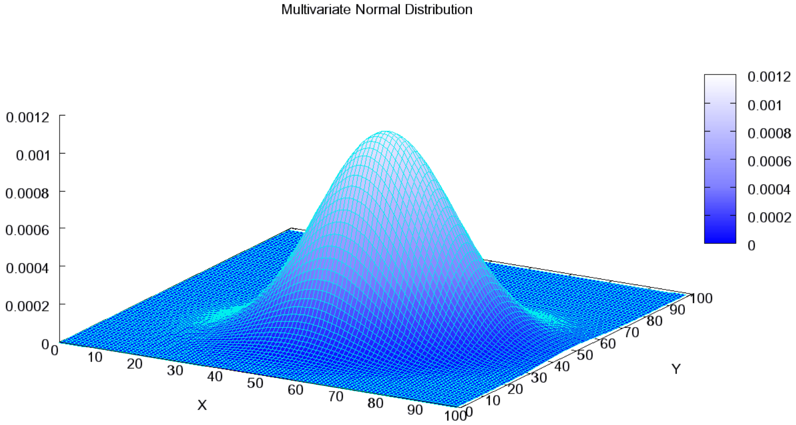

https://commons.wikimedia.org/w/index.php?curid=1260349

Some important definitions

Let (X_1, X_2) be a 2-dimensional continuous random vector with joint density function $f_{12}(x_1, x_2)$. Then the **marginal density function** is defined as

$$f_1(x_1) = \int_{-\infty}^{\infty} f_{12}(x_1, x_2)\, dx_2.$$

and the **marginal distribution function** is defined as

$$F_1(x_1) = P(X_1 \leq x_1, X_2 < \infty) = \int_{-\infty}^{x_1} \int_{-\infty}^{\infty} f_{12}(x_1', x_2)\, dx_1' dx_2.$$

Proposition

Necessary and sufficient conditions for a bounded, nondecreasing, and right-continuous function F on $\mathbb{R}^2$ to be a *bivariate distribution function* *are:*

1. $\lim_{x_j \to -\infty} F(x_1, x_2) = 0$, $j = 1, 2$;
2. $\lim_{(x_1, x_2) \to (\infty, \infty)} F(x_1, x_2) = 1$;
3. (rectangle inequality or 2-increasing fct) *for any* (a_1, a_2), (b_1, b_2) *with* $a_1 < b_1$, $a_2 < b_2$,

$$F(b_1, b_2) - F(a_1, b_2) - F(b_1, a_2) + F(a_1, a_2) \geq 0.$$

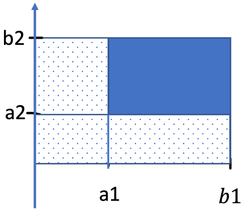

Quantile function

Let X be a real-valued random variable with distribution function $F(x)$. Then the **quantile function** of X is defined as

$$Q(p) = F^{-1}(p) = \inf\{x : F(x) \geq p\}, \quad 0 \leq p \leq 1. \qquad (1)$$

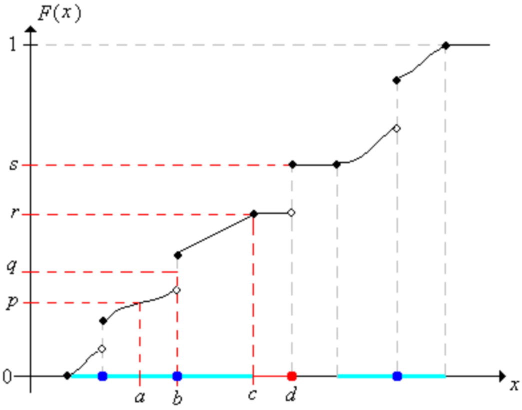

Quantile function

Note: If $F(x)$ is continuous and strictly increasing, $Q(p)$ is the unique value x such that $F(x) = p$.

Exercise 1 Let X be an $\mathrm{Exp}(\lambda)$ random variable:

$$P(X \leq x) = F(x) = 1 - \exp[-\lambda x]$$

for $x \geq 0$ and $\lambda > 0$. Compute the quantile function $Q(p)$.

Solution:

Quantile function

Note: If $F(x)$ is continuous and strictly increasing, $Q(p)$ is the unique value x such that $F(x) = p$.

Exercise 1 Let X be an $\mathrm{Exp}(\lambda)$ random variable:

$$P(X \leq x) = F(x) = 1 - \exp[-\lambda x]$$

for $x \geq 0$ and $\lambda > 0$. Compute the quantile function $Q(p)$.

Solution: $Q(p) = \lambda^{-1}(-\log(1-p))$.

Quantile function

Note: If $F(x)$ is continuous and strictly increasing, $Q(p)$ is the unique value x such that $F(x) = p$.

Exercise 1 Let X be an $\mathrm{Exp}(\lambda)$ random variable:

$$P(X \leq x) = F(x) = 1 - \exp[-\lambda x]$$

for $x \geq 0$ and $\lambda > 0$. Compute the quantile function $Q(p)$.

Solution: $Q(p) = \lambda^{-1}(-\log(1 - p))$.

Exercise 2 Show: if X (continuous) has quantile function $Q(p)$, then $1/X$ has quantile function $1/Q(1 - p)$.

Solution:

Quantile function

Note: If $F(x)$ is continuous and strictly increasing, $Q(p)$ is the unique value x such that $F(x) = p$.

Exercise 1 Let X be an $\mathrm{Exp}(\lambda)$ random variable:

$$P(X \le x) = F(x) = 1 - \exp[-\lambda x]$$

for $x \ge 0$ and $\lambda > 0$. Compute the quantile function $Q(p)$.

Solution: $Q(p) = \lambda^{-1}(-\log(1-p))$.

Exercise 2 Show: if X (continuous) has quantile function $Q(p)$, then $1/X$ has quantile function $1/Q(1-p)$.

Solution:

$1 - p = P(X \ge x) = P(1/X \le 1/x)$; thus, $Q(1-p) = 1/x$.

Uniform distribution on $[a, b]$

For $a, b \in \mathbb{R}$, $a < b$, define the **uniform distribution** on $[a, b]$ by density

$$f(x) = \frac{1}{b-a}.$$

Exercise 3 Find the uniform distribution fct $F(x)$ and its quantile fct.

<u>Solution:</u>

Uniform distribution on $[a, b]$

For $a, b \in \mathbb{R}$, $a < b$, define the **uniform distribution** on $[a, b]$ by density

$$f(x) = \frac{1}{b - a}.$$

Exercise 3 Find the uniform distribution fct $F(x)$ and its quantile fct.

<u>Solution:</u>

$$F(x) = \frac{x - a}{b - a} \quad \text{and} \quad Q(p) = a + (b - a)p.$$

Uniform distribution on $[a, b]$

For $a, b \in \mathbb{R}$, $a < b$, define the **uniform distribution** on $[a, b]$ by density

$$f(x) = \frac{1}{b - a}.$$

Exercise 3 Find the uniform distribution fct $F(x)$ and its quantile fct.

<u>Solution:</u>

$$F(x) = \frac{x - a}{b - a} \quad \text{and} \quad Q(p) = a + (b - a)p.$$

Let U be a uniform random variable on $[0, 1]$; then for $0 \leq c < d \leq 1$,

$$P(c \leq U \leq d) = d - c.$$

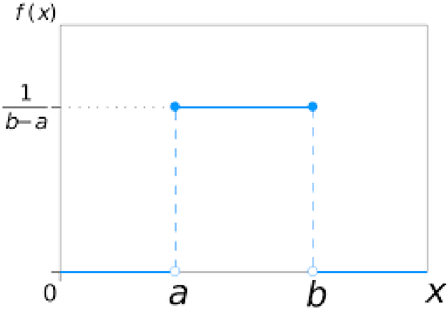
f(x)
1
b−a
0
a
b
x

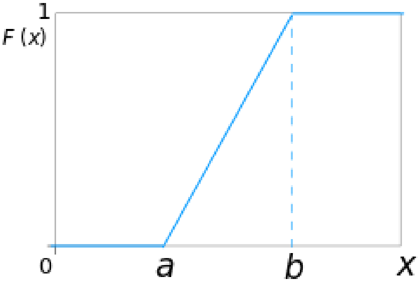
1
F(x)
0
a
b
x

The distribution of random variable $F(X)$

Note: From now on, we only consider **continuous** random variables and random vectors. $X \sim F$ ($\sim$ means "distributed as")

The distribution of random variable $F(X)$

Note: From now on, we only consider **continuous** random variables and random vectors. $X \sim F$ ($\sim$ means "distributed as")

- If X is a random variable with distribution function F_X, then $F_X(X)$ is a transformation of X and again a random variable. What is the distribution of $F_X(X)$?

The distribution of random variable $F(X)$

Note: From now on, we only consider **continuous** random variables and random vectors. $X \sim F$ ($\sim$ means "distributed as")

- If X is a random variable with distribution function F_X, then $F_X(X)$ is a transformation of X and again a random variable. What is the distribution of $F_X(X)$?
- For $0 < u < 1$,

$$P(F_X(X) \leq u) = P(X \leq F_X^{-1}(u)) = F_X(F_X^{-1}(u)) = u$$

- Thus, the distribution of $F_X(X)$ is uniform on $[0, 1]$: $F_X(X) \sim \mathrm{U}(0, 1)$ ("probability integral transform").

The distribution of random variable $F(X)$

Note: From now on, we only consider **continuous** random variables and random vectors. $X \sim F$ ($\sim$ means "distributed as")

- If X is a random variable with distribution function F_X, then $F_X(X)$ is a transformation of X and again a random variable. What is the distribution of $F_X(X)$?
- For $0 < u < 1$,

$$P(F_X(X) \leq u) = P(X \leq F_X^{-1}(u)) = F_X(F_X^{-1}(u)) = u$$

- Thus, the distribution of $F_X(X)$ is uniform on $[0, 1]$: $F_X(X) \sim \mathrm{U}(0, 1)$ ("probability integral transform").
- Let $U \sim \mathrm{U}(0, 1)$ and F any distribution fct. Then $X = F^{-1}(U) \sim F$.
  Proof:

The distribution of random variable $F(X)$

<u>Note:</u> From now on, we only consider **continuous** random variables and random vectors. $X \sim F$ ($\sim$ means "distributed as")

- If X is a random variable with distribution function F_X, then $F_X(X)$ is a transformation of X and again a random variable. What is the distribution of $F_X(X)$?
- For $0 < u < 1$,

$$P(F_X(X) \leq u) = P(X \leq F_X^{-1}(u)) = F_X(F_X^{-1}(u)) = u$$

- Thus, the distribution of $F_X(X)$ is uniform on $[0, 1]$: $F_X(X) \sim \mathrm{U}(0, 1)$ ("probability integral transform").
- Let $U \sim \mathrm{U}(0, 1)$ and F any distribution fct. Then $X = F^{-1}(U) \sim F$.
  Proof: $P(F^{-1}(U) \leq x) = P(U \leq F(x)) = F(x)$, $x \in \mathbb{R}$.
  ("quantile transform")

A Motivating Example (Hofert, 2018)

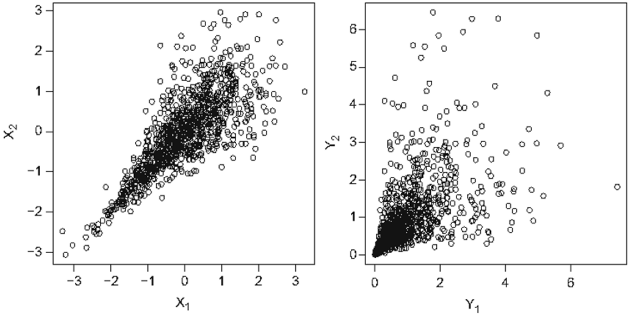

Fig. 1.1 Scatter plots of $n = 1000$ independent observations of (X_1, X_2) (left) and of (Y_1, Y_2) (right)

$$corr(X_1, X_2) = .77 \text{ and } corr(Y_1, Y_2) = .56$$

A Motivating Example (Hofert, 2018)

The marginal distributions of (X_1, X_2) and (Y_1, Y_2)

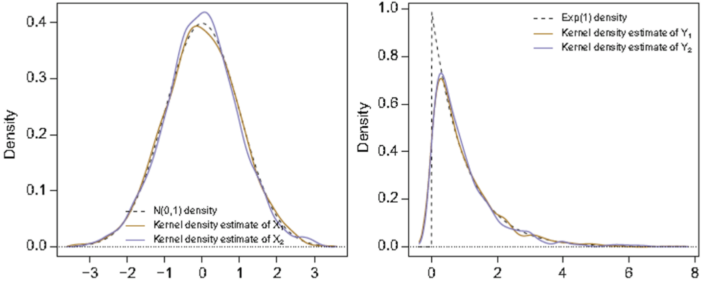

Fig. 1.2 Kernel density estimates of the densities of X_1, X_2 (left) and Y_1, Y_2 (right). The dashed curves represent the $\mathrm{N}(0,1)$ (left) and $\mathrm{Exp}(1)$ (right) densities

If we could transform the two data sets so that they become similar in terms of the underlying marginal dfs, their comparison in terms of dependence would be made on much fairer grounds.

A Motivating Example (Hofert, 2018)

After a probability transformation of the marginals of (X_1, X_2) and $(Y_1, Y_2) \to \; (F_1(X_1), F_2(X_2))$ and $(G_1(Y_1), G_2(Y_2))$

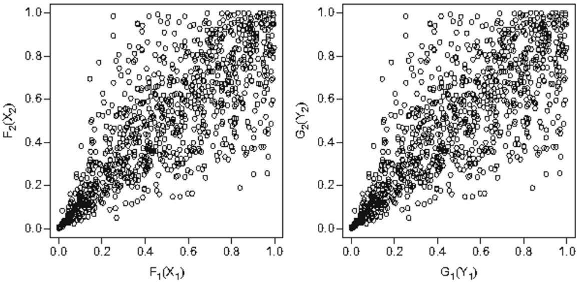

Fig. 1.3 Scatter plots of $n = 1000$ independent observations of the bivariate random vectors $(F_1(X_1), F_2(X_2))$ (left) and $(G_1(Y_1), G_2(Y_2))$ (right)

A Motivating Example (Hofert, 2018)

After a probability transformation of the marginals of (X_1, X_2) and $(Y_1, Y_2) \rightarrow \ (F_1(X_1), F_2(X_2))$ and $(G_1(Y_1), G_2(Y_2))$

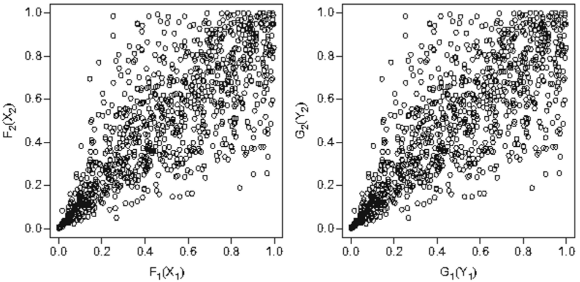

Fig. 1.3 Scatter plots of $n = 1000$ independent observations of the bivariate random vectors $(F_1(X_1), F_2(X_2))$ (left) and $(G_1(Y_1), G_2(Y_2))$ (right)

If dependence between the components of a random vector should not be affected by its marginal distributions, the conclusion is that **the two data sets are indistinguishable in terms of dependence and only differ in terms of the underlying marginal dfs**.

Definition of copula

A copula C is a multivariate distribution with standard uniform margins, that is, $U \sim \mathrm{U}(0,1)$.

Definition of copula

A copula C is a multivariate distribution with standard uniform margins, that is, $U \sim \mathrm{U}(0,1)$.

Characterization in 2-dimensional case:

1. $C : [0,1]^2 \mapsto [0,1]$

Definition of copula

A copula C is a multivariate distribution with standard uniform margins, that is, $U \sim \mathrm{U}(0,1)$.

Characterization in 2-dimensional case:

1. $C : [0,1]^2 \mapsto [0,1]$
2. for every $u, v \in [0,1]$,
   $C(u,0) = 0 = C(0,v)$ and $C(u,1) = u$ and $C(1,v) = v$,

Definition of copula

A copula C is a multivariate distribution with standard uniform margins, that is, $U \sim \mathrm{U}(0,1)$.
Characterization in 2-dimensional case:

1. $C : [0,1]^2 \mapsto [0,1]$
2. for every $u, v \in [0,1]$,
   $C(u,0) = 0 = C(0,v)$ and $C(u,1) = u$ and $C(1,v) = v$,
3. for every $u_1, u_2, v_1, v_2 \in [0,1]$ such that $u_1 \leq u_2$ and $v_1 \leq v_2$,
   $C(u_2,v_2) - C(u_2,v_1) - C(u_1,v_2) + C(u_1,v_1) \geq 0$.

Definition of copula

A copula C is a multivariate distribution with standard uniform margins, that is, $U \sim \mathrm{U}(0,1)$.
Characterization in 2-dimensional case:

1. $C : [0,1]^2 \mapsto [0,1]$
2. for every $u, v \in [0,1]$,
   $C(u,0) = 0 = C(0,v)$ and $C(u,1) = u$ and $C(1,v) = v$,
3. for every $u_1, u_2, v_1, v_2 \in [0,1]$ such that $u_1 \leq u_2$ and $v_1 \leq v_2$,
   $C(u_2,v_2) - C(u_2,v_1) - C(u_1,v_2) + C(u_1,v_1) \geq 0$.

Example: *Independence copula*

$$C(u,v) = uv$$

Check properties 1-3 !

Definition of copula

A copula C is a multivariate distribution with standard uniform margins, that is, $U \sim \mathrm{U}(0,1)$.
Characterization in 2-dimensional case:

1. $C : [0,1]^2 \mapsto [0,1]$
2. for every $u, v \in [0,1]$,
   $C(u,0) = 0 = C(0,v)$ and $C(u,1) = u$ and $C(1,v) = v$,
3. for every $u_1, u_2, v_1, v_2 \in [0,1]$ such that $u_1 \leq u_2$ and $v_1 \leq v_2$,
   $C(u_2,v_2) - C(u_2,v_1) - C(u_1,v_2) + C(u_1,v_1) \geq 0$.

Example: *Independence copula*

$$C(u,v) = uv$$

Check properties 1-3 ! For $u_1 \leq u_2$ and $v_1 \leq v_2$,

$$u_2 v_2 - u_2 v_1 - u_1 v_2 + u_1 v_1 = u_2(v_2 - v_1) - u_1(v_2 - v_1) = (u_2 - u_1)(v_2 - v_1) \geq 0$$

Sklar's Theorem, 1959

(Remember: assuming continuous random variables).
For **any** bivariate distribution function $F(x_1, x_2)$ with margins $F_1(x_1)$ and $F_2(x_2)$ there exists a **unique** copula C such that

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)).$$

Sklar's Theorem, 1959

(Remember: assuming continuous random variables).
For **any** bivariate distribution function $F(x_1, x_2)$ with margins $F_1(x_1)$ and $F_2(x_2)$ there exists a **unique** copula C such that

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)).$$

Moreover, if $F_1(x_1)$ and $F_2(x_2)$ are the marginals of $F(x_1, x_2)$, then copula C can be written as

$$C(u_1, u_2) = F(F_1^{-1}(u_1), F_2^{-1}(u_2)),$$

where F_1^{-1}, F_2^{-1} are the quantile functions of the margins.

Sklar's Theorem, 1959

(Remember: assuming continuous random variables).
For **any** bivariate distribution function $F(x_1, x_2)$ with margins $F_1(x_1)$ and $F_2(x_2)$ there exists a **unique** copula C such that

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)).$$

Moreover, if $F_1(x_1)$ and $F_2(x_2)$ are the marginals of $F(x_1, x_2)$, then copula C can be written as

$$C(u_1, u_2) = F(F_1^{-1}(u_1), F_2^{-1}(u_2)),$$

where F_1^{-1}, F_2^{-1} are the quantile functions of the margins.
A copula "couples" the bivariate distribution with its marginals.

Sklar's Theorem, 1959

(Remember: assuming continuous random variables).
For **any** bivariate distribution function $F(x_1, x_2)$ with margins $F_1(x_1)$ and $F_2(x_2)$ there exists a **unique** copula C such that

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)).$$

Moreover, if $F_1(x_1)$ and $F_2(x_2)$ are the marginals of $F(x_1, x_2)$, then copula C can be written as

$$C(u_1, u_2) = F(F_1^{-1}(u_1), F_2^{-1}(u_2)),$$

where F_1^{-1}, F_2^{-1} are the quantile functions of the margins.
A copula "couples" the bivariate distribution with its marginals.
(i) study the structure of stochastic dependency in a "scale-free" manner, i.e., independent of the specific marginal distributions,

(ii) construct families of multivariate distributions with specified properties.

Copula publications/citations in Web-of-Science 1995-2021

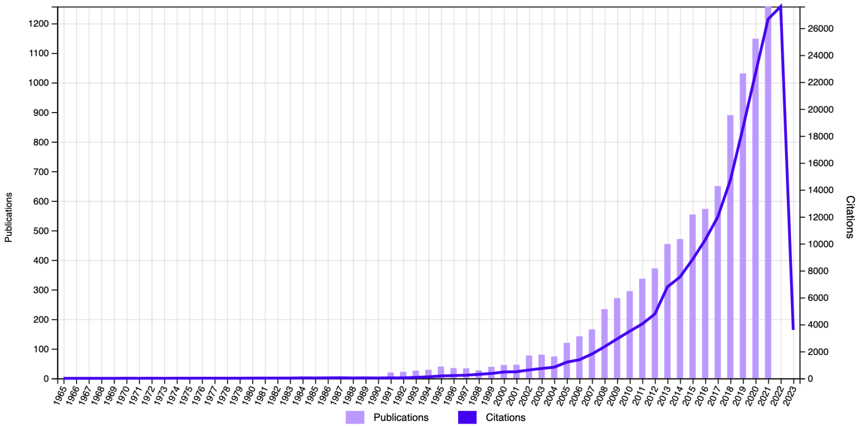

Copula publications/citations in Web-of-Science 1995-2021

Abe Sklar (1925-2020)

Depend. Model. 2021; 9:200–224 DE GRUYTER

Interview Article
Special Issue in memory of Abe Sklar

Open Access

Christian Genest*

A tribute to Abe Sklar

https://doi.org/10.1515/demo-2021-0110
Received June 27, 2021; accepted July 2, 2021

This paper gives an account of the life and works of the American mathematician Abe Sklar. Born in Chicago on November 17, 1925, Sklar completed his PhD at the California Institute of Technology in 1956. He then joined the Illinois Institute of Technology, where he taught mathematics until his retirement in 1995. With his close friend and lifelong collaborator Berthold Schweizer (1929–2010), he was a pioneer of the theory of probabilistic metric spaces, which were introduced in 1942 by the Austro-American mathematician Karl Menger (1902–85). Together, Schweizer and Sklar made important contributions to the algebra of functions, the study of t-norms, and distributional chaos. Sklar is also credited for the notion of copula and for showing that any multivariate distribution function can be expressed in terms of its univariate margins and a copula. This result, known as Sklar's representation theorem, is the bedrock of a widespread data analytical technique called copula modeling. Sklar passed away in Chicago on October 30, 2020.

By now, just about anyone who is conscious of the role of dependence in data analysis has heard of copulas as a powerful and flexible tool for modeling association and assessing its impact on inference, decision making, and risk management. This approach is rooted in a 3-page note, written in French, which appeared in 1959 in the *Publications de l'Institut de statistique de l'Université de Paris*. This paper [S6], attributed only to "M. Sklar" (M. for Mr.), without address or affiliation, claimed without proof that given any d-variate cumulative distribution function H with one-dimensional margins $F_1, \ldots, F_d$, a function $C : [0,1]^d \to [0,1]$ having specific analytical properties can always be found such that

$$H = C(F_1, \ldots, F_d). \quad (1)$$

https://www.wired.com/2009/02/wp-quant/

FELIX SALMON BUSINESS FEB 23, 2009 12:00 PM

Recipe for Disaster: The Formula That Killed Wall Street

In the mid-'80s, Wall Street turned to the quants—brainy financial engineers—to invent new ways to boost profits. Their methods for minting money worked brilliantly... until one of them devastated the global economy.

In the mid-'80s, Wall Street turned to the quants—brainy financial engineers—to invent new ways to boost profits. Their methods for minting money worked brilliantly... until one of them devastated the global economy. © JIM KRANTZ / INDEX STOCK IMAGERY, INC. / GALLERY STOCK

A YEAR AGO, it was hardly unthinkable that a math wizard like David X. Li might

TRENDING NOW

The formula that killed Wall Street

$$\Pr[T_A < 1, T_B < 1] = \phi_2(\phi^{-1}(F_A(1)), \phi^{-1}(F_B(1)), \gamma)$$

The formula that killed so many pension plans: David X. Li's Gaussian copula, as first published in 2000. Investors exploited it as a quick – and fatally flawed – way to assess risk.

Probability

Specifically, this is a joint default probability – the likelihood that any two members of the pool (A and B) will both default. It's what investors are looking for, and the rest of the formula provides the answer.

Survival times

The amount of time between now and when A and B can be expected to default. Li took the idea from a concept in actuarial science that charts what happens to someone's life expectancy when their spouse dies.

Equality

A dangerously precise concept, since it leaves no room for error. Clean equations help both quants and their managers forget that the real world contains a surprising amount of uncertainty, fuzziness, and precariousness.

Copula

This couples (hence the Latinate term copula) the individual probabilities associated with A and B to come up with a single number. Errors here massively increase the risk of the whole equation blowing up.

Distribution functions

The probabilities of how long A and B are likely to survive. Since these are not certainties, they can be dangerous: Small miscalculations may leave you facing much more risk than the formula indicates.

Gamma

The all-powerful correlation parameter, which reduces correlation to a single constant – something that should be highly improbable, if not impossible. This is the magic number that made Li's copula function irresistible.

Felix Salmon (Wired, 2009; reprinted in Significance, 2012)

Some references

- Hofert, M, Kojadinovic, I, Mächler, M., Yan, J. (2018) *Elements of Copula Modeling with R.* Springer-Verlag, UseR! Series
- Durante, F., & Sempi, C. (2015). *Principles of copula theory.* Boca Raton, FL: CRC Press.
- Joe, H. (2014). *Dependence modeling with copulas.* Boca Raton, FL: Chapman & Hall/CRC.
- Jaworski, P., Durante, F., Härdle, W. K., & Rychlik, T. (Eds.) (2010). *Copula theory and itsapplications.* Lecture notes in statistics (Vol. 198). Berlin: Springer
- Sklar, A. (1959). *Fonctions de répartition à n dimensions et leurs marges.* Publications de l'Institut de Statistique de l'Université de Paris, 8, 229–231.
- Nelsen, R. B. (2006). *An introduction to copulas.* New York: Springer.
- Groesser & Okhrin (2020) Wires Computational Statistics: *Copulae: Overview and Recent Developments.*

Some references

Bo Chang

- HOME
- BIO
- PUBLICATIONS
- BLOG

I am a software engineer at Google Brain, based in Toronto, Canada. Prior to that, I was a machine learning researcher at Borealis AI. I finished my Ph.D. in statistics at the University of British Columbia.

Email: contact@bochang.mecontact@bochang.me