Semiparametric Bayesian estimation in copula models

Clara Grazian

Nuffield Department of Medicine, University of Oxford

clara.grazian@ndm.ox.ac.uk

Joint work with Brunero Liseo

BIRS Workshop on Validating and Expanding Approximate Bayesian Computation

BANFF

21 February 2017

There is ABC.....

ABC Rejection algorithm

For $l = 1, \cdots, N$ **do**

    Repeat

        Generate θ' from the prior distribution $\pi(\cdot)$

        Generate z from the likelihood $f(\cdot \mid \theta')$

    Until $\rho(\eta(\mathbf{z}), \eta(\mathbf{y})) < \varepsilon$

    Set $\theta_l = \theta'$

End For

...and ABC

- Bayesian synthetic likelihood [Drovandi et al., 2015, Price et al., 2016]

$$\pi_n(\theta|s_{obs}) \propto \mathcal{N}(s_{obs}; \mu_n(\theta), \Sigma_n(\theta))\pi(\theta)$$

- Bayesian empirical likelihood [Mengersen et al., 2013]

$$\pi(\theta|y) \propto L_{EL}(\theta)\pi(\theta)$$

- Bayesian bootstrap likelihood [Zhu et al., 2015]

$$\pi(\theta|y) \propto L_{BL}(\theta)\pi(\theta)$$

Empirical Likelihood

Empirical likelihood is a way of producing a nonparametric likelihood for a quantity of interest [Owen, 2001]. Schennach [2005] proposes a **Bayesian exponentially tilted empirical likelihood**.

Consider a given set of generalized **moment conditions**

$$E_F\left(h(X,\varphi)\right) = 0,$$

where $h(\cdot)$ is a known function, and φ is the quantity of interest.

$L_{BEL}(\varphi; x)$ is defined as the system of weigths $(p_1, \cdots, p_n)$ obtained as solution of

$$\max_{(p_1,\dots,p_n)} \sum_{i=1}^{n} (-p_i \log p_i)$$

under constraints

$$0 \le p_i \le 1, \qquad \sum_{i=1}^{n} p_i = 1 \qquad \sum_{i=1}^{n} h(x_i,\varphi) p_i = 0$$

[Owen's maximisation problem was $\max_{(p_1,\dots,p_n)} \prod_{i=1}^{n} p_i$.]

The Bayesian use of the empirical likelihood I

We are interested in a function φ and in its posterior

$$\pi(\phi|y) \propto \int_N p(y|\nu,\phi)\pi(\nu|\phi)\pi(\phi)d\nu$$

or

$$\pi(\phi|y) \propto \lim_{N\to\infty} \int_N p(y|\xi_N,\phi)\pi(\xi_N|\phi)\pi(\phi)d\xi_N$$

Then the distribution C can be represented as

$$C = (\varphi, C^*)$$

where C^* belongs to an infinite dimensional metric space (H, d_H).

L_{BEL} may be seen as the derivation of the integrated likelihood for ϕ

$$L_{BEL}^{(\lambda)}(\phi; y) = \int_\Xi L(\phi, \xi; y)d\Pi(\xi)$$

where $\Pi(\xi)$ is the prior process implicitly induced by L_{BEL}.

The Bayesian use of the empirical likelihood II

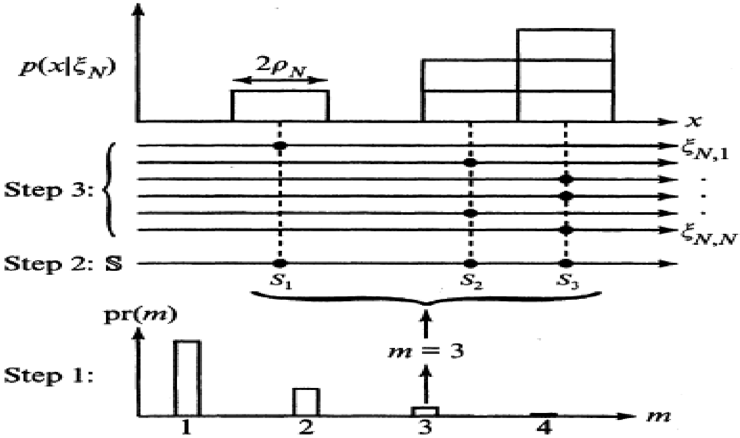

Why copulas?

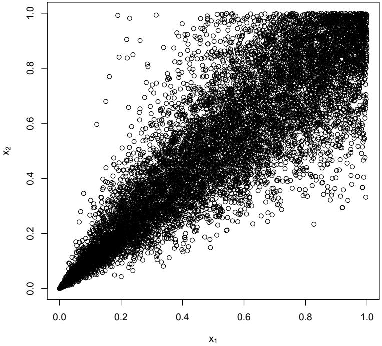

Figure: The Clayton copula exhibits greater dependence in the negative tail

Sklar's Theorem

A copula model is a way of representing the joint distribution of a random vector $\mathbf{X} = (X_1, \ldots, X_d)$. Given an d-variate cumulative distribution function (CDF) $\mathbf{F}$, it is possible to show [Sklar, 1959] that there always exists an d-variate function $C : [0,1]^d \to [0,1]$, such that

$$\mathbf{F}(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d))$$

where F_j is the marginal CDF of X_j.
Therefore, in case that the multivariate distribution has a density $\mathbf{f}$, and this is available, it holds further that

$$\mathbf{f}(x_1, \ldots, x_d) = c(F_1(x_1), \ldots, F_d(x_d)) \cdot f_1(x_1) \cdots f_d(x_d)$$

Example of copula functions

- Clayton copula: $C_\theta(u,v) = [\max\{u^{-\theta} + v^{-\theta} - 1; 0\}]^{-\frac{1}{\theta}}$ for $\theta \in [-1,1) \setminus \{0\}$
- Gumbel copula: $C_\theta(u,v) = \exp[-((-\log(u)^\theta + (-\log(v))^\theta))^{\frac{1}{\theta}}]$ for $\theta \in [1,\infty]$
- Frank copula $C_\theta(u,v) = -\frac{1}{\theta} \log\left[1 + \frac{(\exp(-\theta u)-1)(\exp(-\theta v)-1)}{\exp(-\theta)-1}\right]$ for $\theta \in \mathbb{R} \setminus 0$

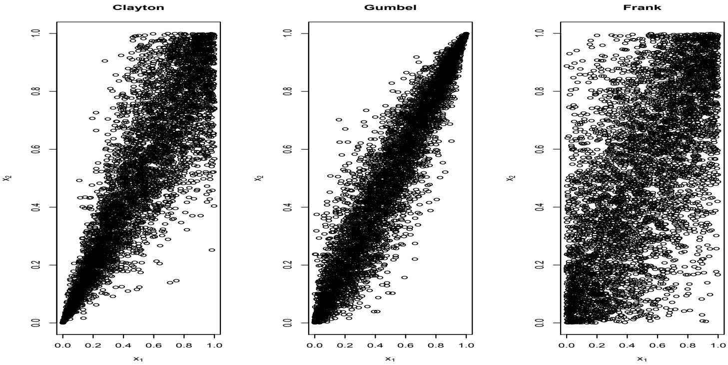

Figure: Simulations from different copula functions

Estimation methods for copula models

- **Frequentist methods:**
  - *Inference from the margins* [Joe, 2015]
  - method of moments [Oh and Patton, 2013]
  - semiparametric approach [Genest et al., 1995]
- **Bayesian methods:** [Smith, 2013]
  - multivariate discrete data [Smith and Khaled, 2012]
  - conditional copulae [Craiu and Sabeti, 2012]
  - vine-copulae [Min and Czado, 2010]
  - nonparametric approach [Wu et al., 2014]

Likelihood function

The likelihood function is complicated!

In the continuous case, the density of a multivariate distribution, in its copula representation is

$$f(x|\lambda,\theta) = c(u;\theta)\prod_{j=1}^{d} f_j(x_j|\lambda_j)$$

where $u = (u_1, \cdots, u_d) = (F_1(x_1;\lambda_1), \cdots, F_d(x_d;\lambda_d))$.

The posterior distribution for (θ,λ) is

$$\pi(\theta,\lambda|x) \propto \pi(\theta,\lambda)\prod_{i=1}^{n}\left[c(u_i;\theta)\prod_{j=1}^{d} f(x_{ij};\lambda_j)\right].$$

Remark: the likelihood function is not separable in $\lambda_1, \cdots, \lambda_d$ and θ because u_i depends on the marginal parameter λ.

A semiparametric approach

Why a semiparametric approach?

- if the interest is in a **functional of the dependence**, the likelihood function for it may be very complicated

  Example: Clayton copula

  - likelihood: $Ł(\theta; u, v) = \prod_{i=1}^{n} (\theta + 1)(u_i v_i)^{-(\theta+1)}(u_i^{-\theta} + v_i^{-\theta} - 1)^{-\frac{2\theta+1}{\theta}}$
  - functionals: $\tau = \frac{\theta}{\theta+2}$ *and* $\lambda_L = 2^{-\frac{1}{\theta}}$ *but* $\rho = \cdots$

- methods of selection of the copula may be unreliable

In this situation we derive an approximated posterior distribution

$$\pi(\phi|x) \propto \pi(\phi) L_{BEL}(\phi; x)$$

ABSCop: STEP 1

ABSCop Step 1: Marginal Estimation

Given a sample $X = (X_1, X_2, \cdots, X_d)$ with joint cdf $F_X(x)$ and marginal cdf's $F_1(x_1; \lambda_1), \cdots, F_d(x_d; \lambda_d)$

For $j = 1, \cdots, d$

> Derive a posterior sample for λ_j: $(\lambda_j^1, \cdots, \lambda_j^{S_j})$ approximating the marginal posterior $\pi(\lambda_j | x_j)$

End For

ABSCop Step 2: Joint Estimation

For $b = 1, \cdots, B$

- Draw $\phi^{(b)} \sim \pi(\phi)$
- Sample one value λ^{s_j} from each marginal posterior sample: $\lambda' = (\lambda_1^{(s_1)}, \cdots, \lambda_d^{(s_d)})$
- Derive a matrix of uniformly distributed pseudo-data $u_{ij} = F_j(x_{ij}; \lambda_j^{(s_j)})$

$$u' = \begin{pmatrix} u_{11}^{(s_1)} & u_{12}^{(s_2)} & \dots & u_{1d}^{(s_d)} \\ u_{21}^{(s_1)} & u_{22}^{(s_2)} & \dots & u_{2d}^{(s_d)} \\ \dots & \dots & u_{ij}^{(s_j)} & \dots \\ u_{n1}^{(s_1)} & u_{n2}^{(s_2)} & \dots & u_{nd}^{(s_d)} \end{pmatrix}.$$

- Compute $L_{BEL}(\varphi^{(b)}; u') = \omega_b$

End For

Validation: nonparametric estimation of the marginals

Suppose $(X_{11}, X_{21}, \cdots, X_{d1}), \cdots, (X_{1n}, X_{2n}, \cdots, X_{dn})$ are independent random vectors with distribution function $\mathbf{F}$ and marginal $F_1, F_2, \cdots, F_d$.

The empirical estimator of the copula function $C(u_1, u_2, \cdots, u_n) = F(F_1^{-1}(u_1), F_2^{-1}(u_2), \cdots, F_d^{-1}(u_d))$ is

$$C_n(u_1, u_2, \cdots, u_d) = \mathbf{F}_n(F_{1n}^{-1}(u_1), F_{2n}^{-1}(u_2), \cdots, F_{dn}^{-1}(u_d)),$$

where $\mathbf{F}_n, F_{1n}, F_{2n}, \cdots, F_{dn}$ are the joint and marginal empirical distribution functions of the observations.

The empirical copula process is defined as

$$\mathbb{C}_n = \sqrt{n}(C_n - C)$$

and if the j-th first order partial derivative exists and is continous on $V_{d,j} = \{u \in [0,1]^d : 0 < u_j < 1\}$, then $\mathbb{C}_n$ converges weakly to the Gaussian process $\{\mathbb{G}_C(u_1, u_2, \cdots, u_d), 0 < u_1, u_2, \cdots, u_d < 1\}$in $\ell^\infty([0,1]^d)$.

ABSCop algorithm

Goal: *estimating a functional of the dependence (Spearman's ρ, Kendall's τ, tail dependence coefficients λ_L and λ_U, etc.)*

- Select a quantity of interest ϕ and a prior $\pi(\phi)$

$$\rho = 12 \int_0^1 \int_0^1 C(u_j, u_h) du_j du_h - 3.$$

  with $\pi(\rho) \sim \mathscr{U}(-1,1)$.

- Select a (nonparametric) estimators ϕ_n

$$\rho_n = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{12}{n^2-1} R_i Q_i \right) - 3 \frac{n+1}{n-1},$$

- Compute the empirical likelihood of ϕ based on its estimate
- Derive via simulation the posterior distribution $\pi(\phi; \mathbf{x})$

Clayton and Frank, $d = 2$

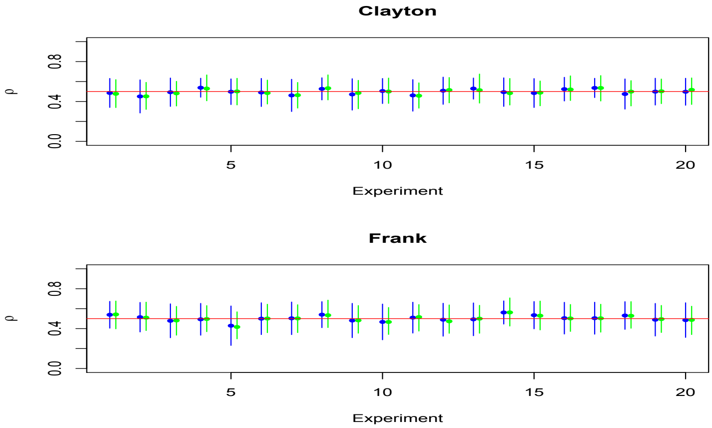

Figure: Comparison between frequentist (blue) and Bayesian estimates (green). 20 out of 500 experiments with simulations from a Clayton copula (above) and a Frank copula (below) ($n = 1000$).

What happens when $\rho \to 1$

Borkowf [2002] shows that the asymptotic variance of ρ_n is

$$\sigma^2(\rho_n) = 144(-9\theta_1^2 + \theta_2 + 2\theta_3 + 2\theta_4 + 2\theta_5), \tag{1}$$

where

$$\begin{aligned}
\theta_1 &= \mathbb{E}[F_1(X_1)F_2(Y_1)] \\
\theta_2 &= \mathbb{E}[(1 - F_1(X_1))^2(1 - F_2(Y_1))^2] \\
\theta_3 &= \mathbb{E}[(1 - F(X_1, Y_2))(1 - F(X_2))(1 - F(Y_1))] \\
\theta_4 &= \mathbb{E}[(1 - F_1(\max\{X_1, X_2\}))(1 - F_2(Y_1))(1 - F_2(Y_2))] \\
\theta_5 &= \mathbb{E}[(1 - F_1(X_1))(1 - F_1(X_2))(1 - F_2(\max\{Y_1, Y_2\}))].
\end{aligned}$$

Consistent estimates of the above quantities are available in Genest & Favre [2007].
However, in the case of **perfect rank agreement**, when plugging-in the sample estimates of the θ_j's into expression (1), one gets a negative number.

Intervals length for $d = 2$

Table: Simulations from different copulas: average length and empirical coverage based of the intervals obtained both via frequentist and Bayesian methods, based on 500 repetitions of the experiment

		Ave. Length	**Coverage**
Clayton ($\rho = 0.5$)	*Freq.*	0.2664	0.998
	Bayes.	0.2597	1.000
Frank ($\rho = 0.5$)	*Freq.*	0.3172	1.000
	Bayes.	0.2735	1.000
Gumbel ($\rho = 0.68$)	*Freq.*	-	-
	Bayes.	0.2966	1.000
Gaussian ($\rho = 0.8$)	*Freq.*	-	-
	Bayes.	0.2931	1.000

Comparison with parametric methods

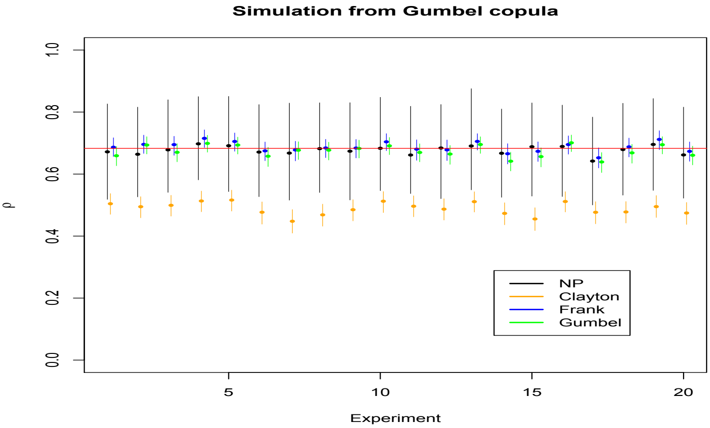

Figure: Bayesian point estimates (points) and credible intervals for 20 out of 500 experiments with data from a Gumbel copula with $\theta = 2$, obtained by specifying a Clayton model (orange), a Frank model (blue) and a Gumbel model (green) or by using our semiparametric approach (black).

Tail dependence I

The **upper and lower tail dependence indices** are defined

$$\lambda_U = \lim_{u\to 1} \Pr\{X_i > F_i^{-1}(u) | X_j > F_j^{-1}(u)\}$$
$$\lambda_L = \lim_{v\to 1} \Pr\{X_i \leq F_i^{-1}(u) | X_j \leq F_j^{-1}(u)\}$$

but may be rewritten in terms of copulas

$$\lambda_U = \lim_{v\to 1} \frac{1-2v-C(v,v)}{1-v}, \qquad \lambda_L = \lim_{v\to 0} \frac{C(v,v)}{v}.$$

that may be estimated by [Joe et al., 1992]

$$\hat{\lambda}_U = 2 - \frac{n}{k}\left\{1-\hat{C}_n\left(\frac{n-k}{n},\frac{n-k}{n}\right)\right\}, \qquad \hat{\lambda}_L = \frac{n}{k}\hat{C}_n\left(\frac{k}{n},\frac{k}{n}\right)$$

Tail dependence II

Schmidt and Stadtmüller [2006] prove

- strong consistency
- asymptotic normality for these estimators.
- derive the asymptotic variance

Tail dependence

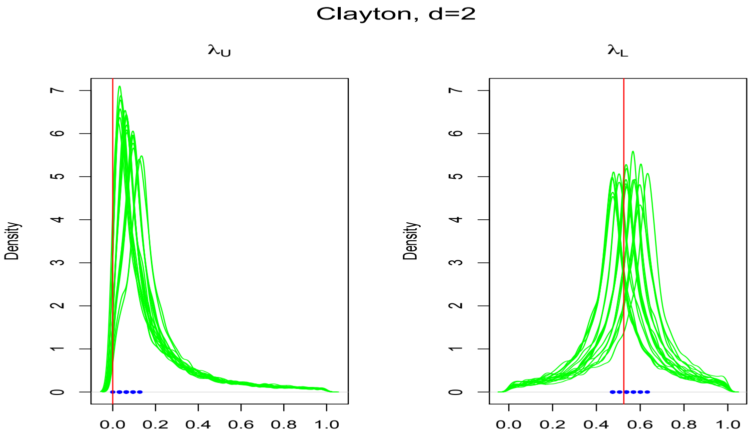

Figure: Comparison between frequentist (blue) and Bayesian (green) estimates for λ_U (left) and λ_L (right). 20 out of 500 simulations from a Clayton copula with $\theta = 1.076$ $(n = 1000)$. The true values are $\lambda_U^{true} = 0$ and $\lambda_L^{true} = 2^{-\frac{1}{\theta}}$ (red lines).

Multivariate Analysis I

Goal: *estimating a functional of the dependence (Spearman's* ρ*, Kendall's* τ*, tail dependence coefficients* λ_L *and* λ_U*, etc.)*

- Select a quantity of interest ϕ and a prior $\pi(\phi)$

$$\rho_1 = \frac{\int_{[0,1]^d} (C(u) - \Pi(u))\, du}{\int_{[0,1]^d} (M(u) - \Pi(u))\, du} = h(d)\left\{2^d \int_{[0,1]^d} C(u)du - 1\right\},$$

where $h(d) = (d+1)/\{2^d - (d+1)\}$ or

$$\rho_2 = h(d)\left\{2^d \int_{[0,1]^d} \Pi(u)dC(u) - 1\right\}.$$

Multivariate Analysis II

- Select a (nonparametric) estimators ϕ_n

$$\hat{\rho}_{1n} = h(d)\left\{2^d \int_{[0,1]^d} \hat{C}_n(u)du - 1\right\} = h(d)\left\{\frac{2^d}{n}\sum_{i=1}^{n}\prod_{j=1}^{d}(1-\hat{U}_{ij}) - 1\right\}$$

$$\hat{\rho}_{2n} = h(d)\left\{2^d \int_{[0,1]^d} \Pi(u)d\hat{C}_n(u) - 1\right\} = h(d)\left\{\frac{2^d}{n}\sum_{i=1}^{n}\prod_{j=1}^{d}\hat{U}_{ij} - 1\right\}.$$

Asymptotic properties of these estimators are explored and assessed in Schmid and Schmidt [2007]. In particular it is known that

$$\sqrt{n}(\hat{\rho}_{kn} - \rho_k) \dot{\sim} \mathcal{N}(0, \sigma_k^2), \qquad k = 1, 2.$$

Multivariate ρ

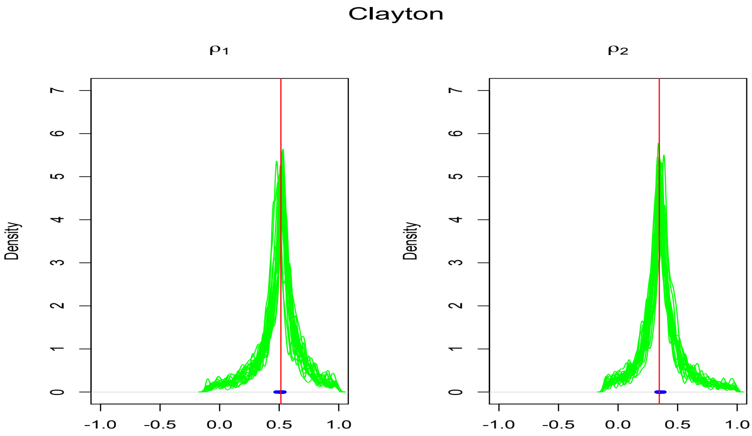

Figure: Comparison between frequentist (blue) and Bayesian (green) estimates of ρ_1 (left) and ρ_2 (right). 20 out of 500 experiments with simulation from a Clayton copula with $\theta = 1.076$ ($n = 1000$).

Dimension $d \geq 2$

Table: Average lengths of the confidence intervals (based on a bootstrap estimator of the variance of the estimates) and of the corresponding Bayesian credible intervals obtained in 50 repetitions of each experiment of dimension d by simulating data from a Clayton copula with $\theta = 1.076$.

	$\hat{\rho}_1^{freq}$	$\hat{\rho}_2^{freq}$	$\hat{\rho}_1^{Bayes}$	$\hat{\rho}_2^{Bayes}$
$d = 2$	0.0032	0.0032	1.1933	1.1801
$d = 3$	0.0026	0.0026	1.0844	1.0853
$d = 4$	0.0026	0.0026	0.9495	0.9594
$d = 5$	0.0027	0.0027	0.8728	0.8914
$d = 6$	0.0027	0.0027	0.8211	0.8224
$d = 7$	0.0030	0.0030	0.8022	0.7882
$d = 8$	0.0031	0.0031	0.7828	0.7541
$d = 9$	0.0032	0.0032	0.7680	0.7492
$d = 10$	0.0035	0.0035	0.7558	0.7439
$d = 25$	0.0047	0.0047	0.7462	0.7480
$d = 50$	0.0073	0.0073	0.7299	0.7634

GARCH(1,1) for Student-t innovation I

Real dataset containing the log-returns FTSE-MIB of five Italian financial institutes

- Monte dei Paschi di Siena
- Banco Popolare
- Unicredit
- Intesa-Sanpaolo
- Mediobanca

by assuming that the log-returns for each bank may be modelled as a **generalized autoregressive conditional heteroscedastic model with parameters** $(1, 1)$ and Student-t innovations.
Data refers to weekdays from 01/07/2013 to 30/06/2014 available on the web-page *https://it.finance.yahoo.com*.

GARCH(1,1) for Student-t innovation II

For $t = 1, \ldots, T$,

$$
\begin{aligned}
y_t &= \varepsilon_t \sqrt{\frac{\nu - 2}{\nu} \omega_t h_t}; \\
h_t &= \alpha_0 + \alpha_1 y_{t-1}^2 + \beta h_{t-1}; \\
\varepsilon_t &\sim \mathcal{N}(0,1); \\
\omega_t &\sim IG\left(\frac{\nu}{2}, \frac{\nu}{2}\right),
\end{aligned}
$$

where $\alpha_0 > 0$, $\alpha_1, \beta >= 0$, $\nu > 2$ and $IG(a,b)$ denotes the invert gamma distribution with shape parameter a and scale parameter b.

Log-returns

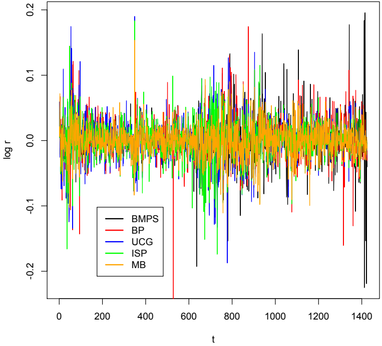

Figure: Log-returns of Monte dei Paschi di Siena (BMPS), Banco Popolare (BP), Unicredit (UCG), Intesa-Sanpaolo (ISP) and Mediobanca (MB) from 01/07/2013 to 30/06/2014, available on the web-page *https://it.finance.yahoo.com*

Posterior Distribution of ρ

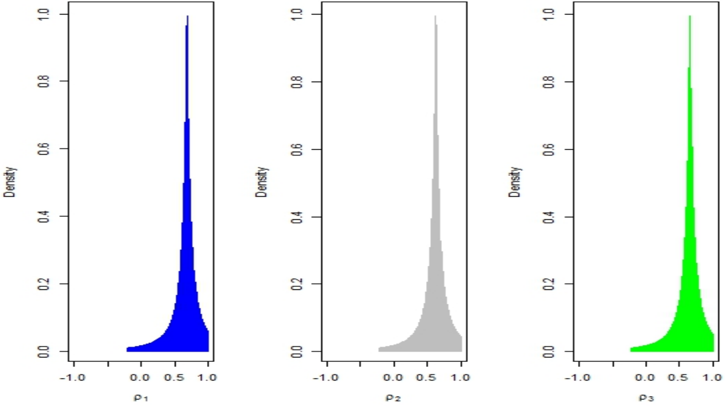

Figure: Approximation of the posterior distribution of the Spearman's ρ for the log-returns of the investments of five Italian institutes based on $10,000$ simulations.

Quantile distributions I

The g-and-k distribution is a popular example of a quantile distribution. This is a transformation of the standard normal distribution function, as follows:

$$Q(z(p);\theta) = a + b\left(1 + c\frac{1-\exp(-gz(p))}{1+\exp(-gz(p))}\right)(1+z(p)^2)^k z(p)$$

where $\theta = (a, b, g, k)$ and c is commonly set fixed at 0.8.

Quantile distributions II

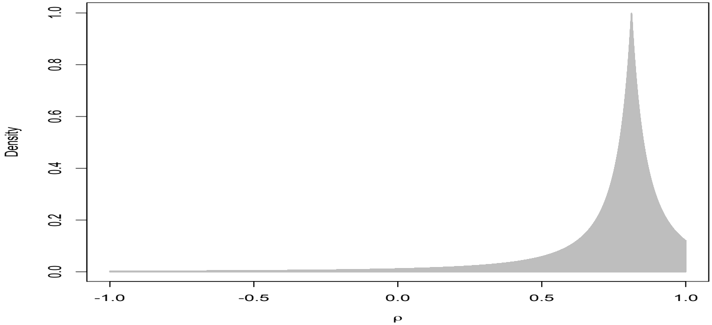

Figure: Spearman's ρ approximated posterior distribution by assuming marginal quantile distributions.

Posterior Distribution of λ_L and λ_U

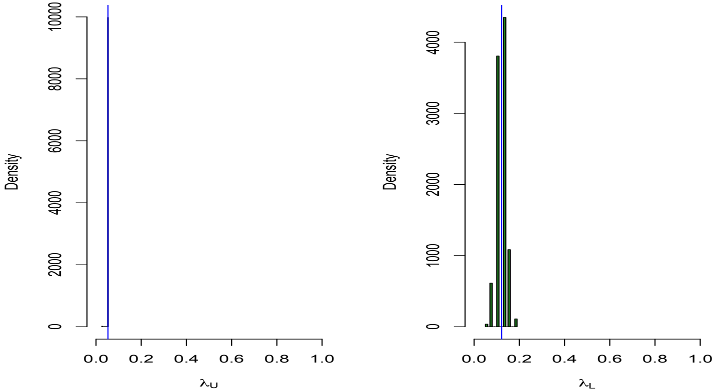

Figure: Approximation of the posterior distribution of the Spearman's ρ for the log-returns of the investments of five Italian institutes based on $10,000$ simulations.

Conditional copulas

A biased estimation of the conditional ρ is

$$\hat{\rho}_n(x) = 12 \sum_{i=1}^{n} w_{ni}(x, h_n)(1 - \hat{U}_{1i})(1 - \hat{U}_{2i}) - 3$$

- bootstrap likelihood [Zhu et al., 2015]
- bootstrap (unbiased) estimator [Lemyre & Quessy, 2016]

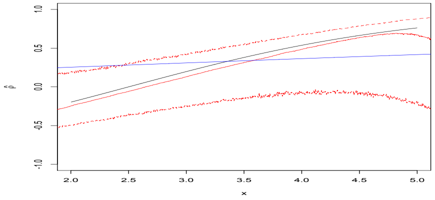

Figure: Simulations from the conditional Clayton copula, true function ρ in black, Bayesian estimates in red, frequentist in blue.

Conclusions

The method presents some advantages and disadvantages.

- ease of elicitation
- robustness in terms of model miss-specification
- generality
- in practical applications, there are often available only asymptotically unbiased estimators
- inefficient with respect to parametric methods (under the assumption that the chosen model is the true one)

Thank you for your attention!

Bibliography I

Borkowf, C. B. (2002) Computing the nonnull asymptotic variance and the asymptotic relative efficiency of Spearma's rank correlation *Computational Statistics & Data Analysis*, 39: 271–286.

Craiu, V. R. and Sabeti, A. (2012) In mixed company: Bayesian inference for bivariate conditional copula models with discrete and continuous outcomes. *J. Multivariate Anal.*, 110: 106–120.

Drovandi, C. C., Pettitt, A. N., and Lee, A. (2015). Bayesian indirect inference using a parametric auxiliary model. *Statistical Science*, 30(1):72–95.

Drovandi, C., Grazian, C., Mengersen, K., and Robert, C. (2017) Approximating the likelihood in Approximate Bayesian computation. *To Appear*.

Ferrari, D., and Zheng, C. (2016) Reliable inference for complex models by discriminative composite likelihood estimation. *ournal of Multivariate Analysis*, 144: 68–80.

Genest, C., and Favre, A. C. (2007). Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of hydrologic engineering*, 12(4), 347-368.

Genest, Christian, Kilani Ghoudi, and L-P. Rivest. (1995) A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika* 82(3): 543-552.

Bibliography II

Grazian, C., and Liseo, B. (2017) Approximate Bayesian Methods for Multivariate and Conditional Copulae. *Advances in Intelligent Systems and Computing*, Volume "Soft Methods for Data Science", vol. 456, pagg. 261–268.

Joe, H. (2005) Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of multivariate analysis*, 94, 401–419.

Joe, H. (2015) *Dependence modeling with copulas*, volume 134 of *Monographs on Statistics and Applied Probability*. CRC Press, Boca Raton, FL.

Joe, Harry and Smith, Richard L. and Weissman, Ishay (1992) Bivariate threshold methods for extremes *Journal of the Royal Statistical Society. Series B.* 54(1): 171–183.

Lemyre, F. C., and Quessy, J. F. (2016) Multiplier bootstrap methods for conditional distributions. *Statistics and Computing*, 1-17.

Lindsay, B.G. (1988) Composite likelihood methods.. *Contemporary mathematics*, 80(1):221–239.

Mengersen, K., Pudlo, P., and Robert, C.P. (2013) Bayesian computation via empirical likelihood. *Proc. of the National Academy of Sciences*, 4(110):1321–1326.

Min, A. and Czado, C. (2010) Bayesian inference for multivariate copulas using pair-copula constructions. *Journal of Financial Econometrics*, 8(4): 511–546.

Bibliography III

Oh, D. H. and Patton, A. J. (2013) Simulated method of moments estimation for copula-based multivariate models. *Journal of the American Statistical Association*, 108(502): 689–700.

Owen, A. B. (2001) *Empirical likelihood*. CRC press.

Price, L. F., Drovandi, C. C., Lee, A., and Nott, D. J. (2016) Bayesian synthetic likelihood. eprints, http://eprints.qut.edu.au/92795/.

Schennach, S. M. (2005) Bayesian exponentially tilted empirical likelihood. *Biometrika*, 92(1): 31–46.

Schmid, F. and Schmidt, R. (2006) Bootstrapping Spearman's multivariate rho. *Proceedings of COMPSTAT 2006*, 759–766.

Schmid, F. and Schmidt, R. (2007) Multivariate Extensions of Spearmans Rho and Related Statistics. *Statistics and Probability Letters*, 77(4):407-416.

Schmid, F. and Stadmüller, U. (2006) Non-parametric estimation of tail dependence. *Scandinavian Journal of Statistics*, 33(2):307-335.

Sklar, M. (1959) Fonctions de répartition à *n* dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231.

Smith, M. S. (2013). Bayesian approaches to copula modelling. In *Bayesian theory and applications*, 336–358. Oxford Univ. Press, Oxford.

Bibliography IV

Smith, M. and Khaled, M. (2012) Estimation of copula models with discrete margins via Bayesian data augmentation. *Journal of the American Statistical Association*, 107(497): 290–303.

Wu, J., Wang, X., and Walker, S. G. (2014). Bayesian nonparametric inference for a multivariate copula function. *Methodology and Computing in Applied Probability*, 16(3): 747-763.

Zhu, W., Marin, J.M., Leisen, F. (2015) A Bootstrap Likelihood approach to Bayesian Computation. *Australian and New Zealand Journal of Statistics*, forthcoming.