Testing hypotheses for the copula of dynamic models

Bruno Remillard, HEC Montréal

BIRS Workshop on Non-Gaussian Multivariate Statistical Models and their Applications

Banff, May 2013

Presentation plan

- Motivation
- Empirical processes
- Goodness-of-fit for copulas
- Example

Motivation

In many financial applications, it is necessary to model both the serial dependance and the dependence between different time series.

This can be done at once by either proposing

- a full parametric model for the multivariate time series

In any case, one has to deal with the residuals of the model since the innovations are not observable.

Motivation

In many financial applications, it is necessary to model both the serial dependance and the dependence between different time series.

This can be done at once by either proposing

- a full parametric model for the multivariate time series
- a hybrid model, i.e.,

In any case, one has to deal with the residuals of the model since the innovations are not observable.

Motivation

In many financial applications, it is necessary to model both the serial dependance and the dependence between different time series.

This can be done at once by either proposing

- a full parametric model for the multivariate time series
- a hybrid model, i.e.,
    - a parametric/semiparametric model for the serial dependance for each series

In any case, one has to deal with the residuals of the model since the innovations are not observable.

Motivation

In many financial applications, it is necessary to model both the serial dependance and the dependence between different time series.

This can be done at once by either proposing

- a full parametric model for the multivariate time series
- a hybrid model, i.e.,
  - a parametric/semiparametric model for the serial dependance for each series
  - a model of the interdependence between the series through copulas.

In any case, one has to deal with the residuals of the model since the innovations are not observable.

To filter or not to filter, that is the question

So far, for copula inference in econometric models, the serial dependence problem is either ignored, i.e., the data are not "filtered" to remove serial dependence, as in Dobrić and Schmid (2005, 2007) and Kole et al. (2007), or the data are "filtered" but the potential inferential problems of using these transformed data are not taken into account.

For example, Panchenko (2005) uses a goodness-of-fit test on "filtered" data (residuals of GARCH models in his case), without proving that his proposed methodology works for residuals. However he mentioned in passing that working with residuals could destroy the asymptotic properties of his test.

A similar situation appears in Breymann et al. (2003) where both the problem of working with residuals and the problem of the estimation of the copula parameters are ignored.

Working with residuals

It complicates the inference procedure since the limiting distribution of statistics and parameters depend in general on unknown parameters (Bai, 1994, Ghoudi and Rémillard, 2004).

In particular, as shown in Bai (2003) and Horváth et al. (2004), the distribution of the empirical process of GARCH residuals in the univariate case (or their squares) is not trivial.

Here, one wants to know what happens in the multivariate case, specially to the limiting process of the empirical copula.

Literature review

The first paper addressing rigorously the problems raised by the use of residuals in estimation and goodness-of-fit of copulas seems to be Chen and Fan (2006).

Using a multivariate GARCH-like model for each univariate time series, they showed the remarkable result that estimating the copula parameters using the rank-based maximum pseudo-likelihood method (Genest et al., 1995) with the ranks of the residuals leads to the same asymptotic distribution as if working with the ranks of innovations.

In particular, the limiting distribution of the estimation of the copula parameters does not depend on the unknown parameters used to estimate the conditional mean and the conditional variance.

This property is crucial if one wants to develop goodness-of-fit tests for the copula family of the innovations.

Multivariate Time Series Model

Stochastic volatility model:

$$\mathbf{X}_i = \boldsymbol{\mu}_i(\boldsymbol{\theta}) + \boldsymbol{\sigma}_i(\boldsymbol{\theta})\boldsymbol{\varepsilon}_i,$$

where the innovations $\boldsymbol{\varepsilon}_i = (\varepsilon_{1i}, \ldots, \varepsilon_{di})^\top$ are i.i.d. with continuous DF K, and $\boldsymbol{\mu}_i$, $\boldsymbol{\sigma}_i$ are $\mathcal{F}_{i-i}$-measurable and independent of $\boldsymbol{\varepsilon}_i$.

Here $\mathcal{F}_{i-1}$ contains information from the past and possibly information from exogenous variables as well.

Since K is continuous, there exists a unique copula C (Sklar, 1959) so that for all $\mathbf{x} = (x_1, \ldots, x_d)^\top \in \mathbb{R}^d$,

$$K(\mathbf{x}) = C\{\mathbf{F}(\mathbf{x})\}, \quad \mathbf{F}(\mathbf{x}) = (F_1(x_1), \ldots, F_d(x_d))^\top,$$

where $F_1, \ldots, F_d$ are the marginal distribution functions of K, i.e., F_j is the distribution function of ε_{ji}, $j = 1, \ldots, d$.

Empirical process of residuals

Given an estimator $\boldsymbol{\theta}_n$ of $\boldsymbol{\theta}$, compute the residuals $\mathbf{e}_{i,n} = (e_{1i,n}, \ldots, e_{di,n})^\top$, where

$$\mathbf{e}_{i,n} = \boldsymbol{\sigma}_i^{-1}(\boldsymbol{\theta}_n)\{\mathbf{X}_i - \boldsymbol{\mu}_i(\boldsymbol{\theta}_n)\}.$$

Further set $\boldsymbol{\Theta}_n = n^{1/2}(\boldsymbol{\theta}_n - \boldsymbol{\theta})$. The main results are deduced from the asymptotic behavior of the process

$$\mathbb{K}_n(s, \mathbf{x}) = \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor ns \rfloor} \{\mathbf{1}(e_{i,n} \leq \mathbf{x}) - K(\mathbf{x})\}, \quad (s, \mathbf{x}) \in [0,1] \times \bar{\mathbb{R}}^d.$$

Here $\mathbf{1}$ stands for the indicator function and $\mathbf{y} \leq \mathbf{x}$ means that the inequality holds componentwise.

Further set

$$K_n(\mathbf{x}) = \frac{1}{n}\sum_{i=1}^{n}\mathbf{1}(e_{i,n} \leq \mathbf{x}), \quad \mathbf{x} \in \bar{\mathbb{R}}^d,$$

and $\mathbf{F}_n(\mathbf{x}) = (F_{1n}(1, x_1), \ldots, F_{dn}(1, x_d))^\top$, where

$$F_{jn}(s, x_j) = \frac{1}{n+1}\sum_{i=1}^{\lfloor ns \rfloor}\mathbf{1}(e_{ji,n} \leq x_j), \quad j = 1, \ldots, d.$$

Weak convergence

From now on, convergence of processes means convergence with respect to the Skorohod topology for the space of càdlàg processes, and is denoted by $\rightsquigarrow$.

To be able to state the convergence result for $\mathbb{K}_n$, one needs to introduce auxiliary empirical processes. Set

$$\alpha_n(s,\mathbf{x}) = \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor ns \rfloor} \{\mathbf{1}(\varepsilon_i \le \mathbf{x}) - K(\mathbf{x})\}, \quad (s,\mathbf{x}) \in [0,1]\times\bar{\mathbb{R}}^d,$$

$$\beta_n(s,\mathbf{u}) = \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor ns \rfloor} \{\mathbf{1}(\mathbf{U}_i \le \mathbf{u}) - C(\mathbf{u})\}, \quad (s,\mathbf{u}) \in [0,1]^{1+d},$$

and $\beta_{j,n}(s,u_j) = \beta_n(s,1,\ldots,1,u_j,1,\ldots,1)$, and $\mathbf{U}_i = \mathbf{F}(\varepsilon_i)$.

Weak convergence to Kiefer processes

It is well known (Bickel and Wichura, 1971) that $\alpha_n \rightsquigarrow \alpha$ and $\beta_n \rightsquigarrow \beta$ where α is a K-Kiefer process and β is a C-Kiefer process.

Recall that α is a K-Kiefer process if it is a continuous centered Gaussian process with $\mathrm{Cov}\,\{\alpha(s,\mathbf{x}),\alpha(t,\mathbf{y})\} = (s\wedge t)\,\{K(\mathbf{x}\wedge\mathbf{y}) - K(\mathbf{x})K(\mathbf{y})\}$, $s\in[0,1]$ and $\mathbf{x},\mathbf{y}\in\mathbb{R}^d$.

Here $(\mathbf{x}\wedge\mathbf{y})_j = \min(x_j, y_j)$, $j = 1\ldots,d$.

Note that for all $(s,\mathbf{x})\in[0,1]\times\bar{\mathbb{R}}^d$, $\alpha(s,\mathbf{x}) = \beta\{s,\mathbf{F}(\mathbf{x})\}$.

Main convergence assumptions

Assume that $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$ are continuously differentiable with respect to $\boldsymbol{\theta} \in \mathcal{O}$, and set $\gamma_{0i} = \boldsymbol{\sigma}_i^{-1} \dot{\boldsymbol{\mu}}_i$ and $\gamma_{1ki} = \boldsymbol{\sigma}_i^{-1} \dot{\boldsymbol{\sigma}}_{ki}$, where $(\dot{\boldsymbol{\mu}}_i)_{jl} = \partial_{\theta_l} \mu_{ji}$, $(\dot{\boldsymbol{\sigma}}_{ki})_{jl} = \partial_{\theta_l} \sigma_{jki} = \partial_{\theta_l} (\boldsymbol{\sigma}_i)_{jk}$.

(A1) $\boldsymbol{\Gamma}_{0,n}(s) = \frac{1}{n} \sum_{i=1}^{\lfloor ns \rfloor} \gamma_{0i}$ and $\boldsymbol{\Gamma}_{1k,n}(s) = \frac{1}{n} \sum_{i=1}^{\lfloor ns \rfloor} \gamma_{1ki}$ converge in prob. to $s\boldsymbol{\Gamma}_0$ and $s\boldsymbol{\Gamma}_{1k}$, uniformly in $s \in [0, 1]$.

(A5) $(\alpha_n, \boldsymbol{\Theta}_n) \rightsquigarrow (\alpha, \boldsymbol{\Theta})$.

Theorem

Under assumptions (A1)–(A7), $\mathbb{K}_n \rightsquigarrow \mathbb{K}$, *with*

$$\mathbb{K}(s,\mathbf{x}) = \alpha(s,\mathbf{x}) + s\nabla K(\mathbf{x})\boldsymbol{\Gamma}_0\Theta + s\sum_{j=1}^{d}\sum_{k=1}^{d} G_{jk}(\mathbf{x})(\boldsymbol{\Gamma}_{1k}\Theta)_j,$$

where $G_{jk}(\mathbf{x}) = f_j(x_j)E\left\{\varepsilon_{k1}\mathbf{1}(\varepsilon_1 \le \mathbf{x})|\varepsilon_{j1} = x_j\right\}$.
Furthermore $\mathbb{F}_{j,n} \rightsquigarrow \mathbb{F}_j$, *where*

$$\begin{aligned}\mathbb{F}_j(s,x_j) &= \beta_j\{s,F_j(x_j)\} + sf_j(x_j)\left\{(\boldsymbol{\Gamma}_0\Theta)_j + x_j(\boldsymbol{\Gamma}_{1j}\Theta)_j\right\} \\ &\quad + s\sum_{k\neq j} f_j(x_j)E(\varepsilon_{k1}|\varepsilon_{j1} = x_j)(\boldsymbol{\Gamma}_{1k}\Theta)_j.\end{aligned}$$

If $\boldsymbol{\sigma}$ *is diagonal, (A7) is not needed for the convergence of* $\mathbb{K}_n$. *In this case,*

$$\mathbb{K}(s,\mathbf{x}) = \alpha(s,\mathbf{x}) + s\nabla K(\mathbf{x})\boldsymbol{\Gamma}_0\Theta + s\sum_{j=1}^{d} G_{jj}(\mathbf{x})(\boldsymbol{\Gamma}_{1j}\Theta)_j.$$

Tests

An immediate application of Theorem 1 if for GOF tests.

GOF tests could also be based on the Rosenblatt transform of K. See, e.g., Genest and Rémillard (2008) and Rémillard (2011) for details.

One can also perform change-point tests.

Define, for all $(s, \mathbf{x}) \in [0,1] \times \bar{\mathbb{R}}^d$, the sequential process

$$\mathbb{A}_n(s, \mathbf{x}) = \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor ns \rfloor} \left\{ \mathbf{1}(e_{i,n} \leq \mathbf{x}) - K_n(\mathbf{x}) \right\}.$$

Many test statistics for detecting structural changes in the innovations are based on $\mathbb{A}_n$. From Theorem 1, one obtains a surprising result:

Corollary

Under assumptions (A1)–(A7), $\mathbb{A}_n \rightsquigarrow \mathbb{A}$, *with*

$$\mathbb{A}(s, \mathbf{x}) = \alpha(s, \mathbf{x}) - s\alpha(1, \mathbf{x}), \quad (s, \mathbf{x}) \in [0, 1] \times \bar{\mathbb{R}}^d.$$

In particular $\mathbb{A}$ *is parameter free, depending only on* K.

Remark

Although the distribution of $\mathbb{A}$ depends on the unknown DF K, it is still possible to bootstrap $\mathbb{A}$, i.e., to generate asymptotically independent copies of $\mathbb{A}$.

Thus it is possible to detect structural changes in the distribution of the innovations using $\mathbb{A}_n$.

The way to do it is to use multipliers. See Rémillard (2012) for details.

Dependence between innovations

Let C be the (unique) copula associated with K.

Since the copula is independent of the margins, one way to estimate it is to remove their effect by replacing $\mathbf{e}_{i,n}$ with the associated rank vectors

$$\mathbf{U}_{i,n} = (U_{1i,n}, \ldots, U_{di,n})^\top, \quad U_{ji,n} = \mathrm{Rank}(e_{ji,n})/(n+1),$$

where $\mathrm{Rank}(e_{ji,n})$ being the rank of $e_{ji,n}$ amongst $e_{j1,n}, \ldots, e_{j1,n}$, $j = 1, \ldots, d$. Also, $\mathbf{U}_{i,n} = \mathbf{F}_n(\mathbf{e}_{i,n})$.

Empirical copula process

Now define the empirical copula

$$C_n(\mathbf{u}) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}(\mathbf{U}_{i,n} \le \mathbf{u}), \quad \mathbf{u} \in [0,1]^d,$$

together with the sequential copula process

$$\mathbb{C}_n(s,\mathbf{u}) = \frac{1}{\sqrt{n}}\sum_{i=1}^{\lfloor ns \rfloor} \left\{\mathbf{1}(U_{i,n} \le \mathbf{u}) - C(\mathbf{u})\right\}, \quad (s,\mathbf{u}) \in [0,1]^{1+d},$$

and set

$$\mathbb{G}_n(s,\mathbf{u}) = \frac{1}{\sqrt{n}}\sum_{i=1}^{\lfloor ns \rfloor} \left\{\mathbf{1}(U_{i,n} \le \mathbf{u}) - C_n(\mathbf{u})\right\}.$$

Main result

Corollary

Under assumptions (A1)–(A7), $\mathbb{C}_n \rightsquigarrow \mathbb{C}$*, with*

$$\mathbb{C}(s,\mathbf{u}) = \check{\mathbb{C}}(s,\mathbf{u}) + s\sum_{j\neq k} \tilde{G}_{jk}(\mathbf{u})(\boldsymbol{\Gamma}_{1k}\Theta)_j,$$

with $\tilde{G}_{jk}$ *deterministic and*

$$\check{\mathbb{C}}(s,\mathbf{u}) = \beta(s,\mathbf{u}) - s\sum_{j=1}^{d} \partial_{u_j} C(\mathbf{u})\beta_j(1,u_j), \quad (s,\mathbf{u}) \in [0,1]^{1+d}.$$

Moreover, $\mathbb{G}_n \rightsquigarrow \mathbb{G}$*, where*

$$\mathbb{G}(s,\mathbf{u}) = \beta(s,\mathbf{u}) - s\beta(1,\mathbf{u}), \quad (s,u) \in [0,1]^{1+d}.$$

Furthermore, under assumptions (A1)–(A6), if the volatility matrices $\boldsymbol{\sigma}_i$ *are diagonal, then* $\mathbb{C}_n \rightsquigarrow \check{\mathbb{C}}$*.*

- An immediate application of Corollary 2 shows that tests for detecting structural change in the copula of the innovations can be based on the process $\mathbb{G}_n$ and that the limiting process $\mathbb{G}$ is parameter free, depending only on the unknown copula C.

- An immediate application of Corollary 2 shows that tests for detecting structural change in the copula of the innovations can be based on the process $\mathbb{G}_n$ and that the limiting process $\mathbb{G}$ is parameter free, depending only on the unknown copula C.
- However, as it was also true for $\mathbb{A}$, it is easy to simulate asymptotically independent copies of $\mathbb{G}$. See Rémillard (2012).

- An immediate application of Corollary 2 shows that tests for detecting structural change in the copula of the innovations can be based on the process $\mathbb{G}_n$ and that the limiting process $\mathbb{G}$ is parameter free, depending only on the unknown copula C.
- However, as it was also true for $\mathbb{A}$, it is easy to simulate asymptotically independent copies of $\mathbb{G}$. See Rémillard (2012).
- The corollary was also used by Duchesne et al. (2012) to build tests of independence between the innovations of several time series.

- An immediate application of Corollary 2 shows that tests for detecting structural change in the copula of the innovations can be based on the process $\mathbb{G}_n$ and that the limiting process $\mathbb{G}$ is parameter free, depending only on the unknown copula C.
- However, as it was also true for $\mathbb{A}$, it is easy to simulate asymptotically independent copies of $\mathbb{G}$. See Rémillard (2012).
- The corollary was also used by Duchesne et al. (2012) to build tests of independence between the innovations of several time series.
- It is remarkable that when the volatility matrices $\boldsymbol{\sigma}_i$ are diagonal, then $\mathbb{C}_n$ converges to $\check{\mathbb{C}}$, which does not depend on $\boldsymbol{\Theta}$, even if $\mathbb{K}$ does.

Example of application

Chen and Fan (2006) studied the dependence of the innovations for the Deutsche Mark/US and Japanese Yen/US exchanges rates, from April 28, 1988 to Dec 31, 1998. AR(3)-GARCH(1,1) and AR(1)-GARCH(1,1) models were fitted on the 2684 log-returns.

Because the series are so long, univariate change-point tests were performed on the standardized residuals and the null hypothesis was accepted.

Then, a copula change-point test was performed ($\mathrm{P-value} = 33\%$, using $N = 100$ replications).

It required 30 hours of calculations, using the multipliers methodology described next.

Figure 1: Residuals vs time.

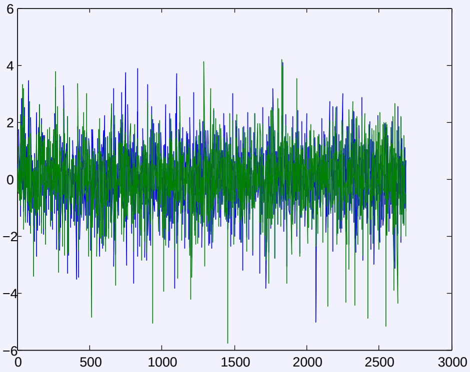

Figure 2: Scatter plot of pseudo-observations.

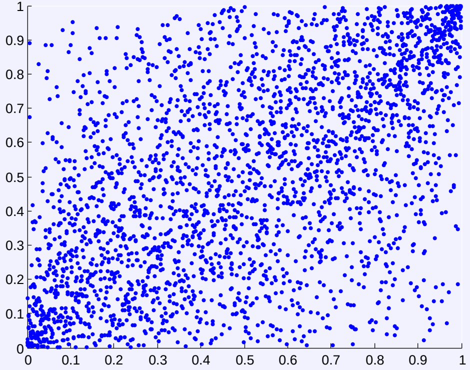

"Best" omnibus test for GOF of copulas

Instead of using $\mathbb{C}_n$ for testing $H_0 : C \in \{C_\theta; \theta \in \mathcal{O}\}$, define pseudo-observations $E_{1,n} = \mathcal{R}_{\theta_n}(e_{1,n})$, ..., $E_{n,n} = \mathcal{R}_{\theta_n}(e_{n,n})$, where R_θ is the Rosenblatt transform of C_θ.

Under the null hypothesis H_0, the empirical distribution function

$$D_n(u) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\left(E_{i,n} \le u\right), \quad u \in [0,1]^d$$

associated with the pseudo-observations $E_1, \ldots, E_n$ should be "close" to the independence copula $C_\perp$.
According to Genest et al. (2009), the best omnibus test for goodness-of-fit is based on

$$S_n^{(B)} = n \int_{[0,1]^d} \{D_n(u) - C_\perp(u)\}^2 \, du.$$

P-values are calculated using parametric bootstrap.

The search for the copula

- The usual copula models considered by Chen and Fan (2006) (Gaussian, Student, Clayton, Frank, Gumbel) were all rejected using $S_n^{(B)}$, while they selected the Student copula as the best model, based on the likelihood rankings.

The search for the copula

- The usual copula models considered by Chen and Fan (2006) (Gaussian, Student, Clayton, Frank, Gumbel) were all rejected using $S_n^{(B)}$, while they selected the Student copula as the best model, based on the likelihood rankings.
- What is the model then?

The search for the copula

- The usual copula models considered by Chen and Fan (2006) (Gaussian, Student, Clayton, Frank, Gumbel) were all rejected using $S_n^{(B)}$, while they selected the Student copula as the best model, based on the likelihood rankings.
- What is the model then?
- The next best model would be a mixture of two Gaussian copulas (Dias and Embrechts, 2004).

The search for the copula

- The usual copula models considered by Chen and Fan (2006) (Gaussian, Student, Clayton, Frank, Gumbel) were all rejected using $S_n^{(B)}$, while they selected the Student copula as the best model, based on the likelihood rankings.
- What is the model then?
- The next best model would be a mixture of two Gaussian copulas (Dias and Embrechts, 2004).
- H_0 was accepted with a 84% p-value, calculated from $N = 100$ replications. The parameters of the two Gaussian copulas are $\hat{\rho} = [0.8205, 0.3749]$ and $\hat{\pi} = [0.4017, 0.5983]$.

Conclusion I

- Parametric bootstrap is a powerful method that works fine but can be quite slow.

Conclusion I

- Parametric bootstrap is a powerful method that works fine but can be quite slow.
- The method is limited by available algorithms for calculating the DF under H_0.

Conclusion I

- Parametric bootstrap is a powerful method that works fine but can be quite slow.
- The method is limited by available algorithms for calculating the DF under H_0.
- One could use Monte Carlo implementation! Replace the DF by an empirical one, obtained from independent Monte Carlo sampling. It is called two-level parametric bootstrap and it works but it is very very slow!

Conclusion I

- Parametric bootstrap is a powerful method that works fine but can be quite slow.
- The method is limited by available algorithms for calculating the DF under H_0.
- One could use Monte Carlo implementation! Replace the DF by an empirical one, obtained from independent Monte Carlo sampling. It is called two-level parametric bootstrap and it works but it is very very slow!
- No such computational problems for tests based on the Rosenblatt transform which is in addition almost always the best!

References I

Bai, J. (1994). Weak convergence of the sequential empirical processes of residuals in ARMA models. *Ann. Statist.*, 22:2051–2061.

Bai, J. (2003). Testing parametric conditional distributions of dynamic models. *The Review of Economics and Statistics*, 85(3):531–549.

Bickel, P. J. and Wichura, M. J. (1971). Convergence criteria for multiparameter stochastic processes and some applications. *Ann. Math. Statist.*, 42:1656–1670.

Breymann, W., Dias, A., and Embrechts, P. (2003). Dependence structures for multivariate high-frequency data in finance. *Quant. Finance*, 3:1–14.

Chen, X. and Fan, Y. (2006). Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification. *Journal of Econometrics*, 135(1-2):125 – 154.

References II

Dias, A. and Embrechts, P. (2004). Dynamic copula models for multivariate high-frequency data in finance. Technical report, ETH Zürich.

Dobrić, J. and Schmid, F. (2005). Testing goodness of fit for parametric families of copulas: Application to financial data. *Comm. Statist. Simulation Comput.*, 34:1053–1068.

Dobrić, J. and Schmid, F. (2007). A goodness of fit test for copulas based on Rosenblatt's transformation. *Comput. Statist. Data Anal.*, 51:4633–4642.

Duchesne, P., Ghoudi, K., and Rémillard, B. (2012). On testing for independence between the innovations of several time series. *Canad. J. Statist.*, 40:447–479.

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

References III

Genest, C. and Rémillard, B. (2008). Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models. *Ann. Inst. H. Poincaré Sect. B*, 44:1096–1127.

Genest, C., Rémillard, B., and Beaudoin, D. (2009). Omnibus goodness-of-fit tests for copulas: A review and a power study. *Insurance Math. Econom.*, 44:199–213.

Ghoudi, K. and Rémillard, B. (2004). Empirical processes based on pseudo-observations. II. The multivariate case. In *Asymptotic Methods in Stochastics*, volume 44 of *Fields Inst. Commun.*, pages 381–406. Amer. Math. Soc., Providence, RI.

Horváth, L., Kokoszka, P., and Teyssière, G. (2004). Bootstrap misspecification tests for ARCH based on the empirical process of squared residuals. *Journal of Statistical Computation and Simulation*, 74(7):469–485.

References IV

Kole, E., Koedijk, K., and Verbeek, M. (2007). Selecting copulas for risk management. *Journal of Banking & Finance*, 31(8):2405 – 2423.

Panchenko, V. (2005). Goodness-of-fit test for copulas. *Phys. A*, 355:176–182.

Rémillard, B. (2011). Validity of the parametric bootstrap for goodness-of-fit testing in dynamic models. Technical report, SSRN Working Paper Series No. 1966476.

Rémillard, B. (2012). Non-parametric change point problems using multipliers. Technical report, SSRN Working Paper Series No. 2043632.

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