



**Probit Transformation
for Nonparametric Kernel Estimation
of the Copula Density**

Gery Geenens
University of New South Wales

Arthur Charpentier
Université du Québec

Davy Paindaveine
SBS-EM, ECARES, Université Libre de Bruxelles

ECARES working paper 2014-23

Probit transformation for nonparametric kernel estimation of the copula density

GERY GEENENS*

School of Mathematics and Statistics,
University of New South Wales, Sydney, Australia,

ARTHUR CHARPENTIER

and DAVY PAINDAVEINE

Département de mathématiques,

ECARES,

Université du Québec à Montréal, Canada,

Université Libre de Bruxelles, Belgium

April 18, 2014

Abstract

Copula modelling has become ubiquitous in modern statistics. Here, the problem of nonparametrically estimating a copula density is addressed. Arguably the most popular nonparametric density estimator, the kernel estimator is not suitable for the unit-square-supported copula densities, mainly because it is heavily affected by boundary bias issues. In addition, most common copulas admit unbounded densities, and kernel methods are not consistent in that case. In this paper, a kernel-type copula density estimator is proposed. It is based on the idea of transforming the uniform marginals of the copula density into normal distributions via the probit function, estimating the density in the transformed domain, which can be accomplished without boundary problems, and obtaining an estimate of the copula density through back-transformation. Although natural, a raw application of this procedure was, however, seen not to perform very well in the earlier literature. Here, it is shown that, if combined with local likelihood density estimation methods, the idea yields very good and easy to implement estimators, fixing boundary issues in a natural way and able to cope with unbounded copula densities. The asymptotic properties of the suggested estimators are derived, and a practical way of selecting the crucially important smoothing parameters is devised. Finally, extensive simulation studies and a real data analysis evidence their excellent performance compared to their main competitors.

Keywords: copula density; transformation kernel density estimator; boundary bias; unbounded density; local likelihood density estimation.

*Corresponding author: ggeenens@unsw.edu.au, School of Mathematics and Statistics, University of New South Wales, Sydney, NSW 2052 (Australia), tel +61 2 938 57032, fax +61 2 9385 7123

1 Introduction

For the last two decades copula modelling has emerged as a major research area of statistics. By definition, a bivariate copula function C is the joint cumulative distribution function (often abbreviated to ‘cdf’ below) of a bivariate random vector whose marginals are Uniform over $[0, 1]$, i.e.,

$$C : \mathcal{I} \doteq [0, 1]^2 \rightarrow [0, 1] : (u, v) \rightarrow C(u, v) = \mathbb{P}(U \leq u, V \leq v),$$

where $U \sim \mathcal{U}_{[0,1]}$, $V \sim \mathcal{U}_{[0,1]}$. Copulas arise naturally in statistics and probability as a mere consequence of two well-known facts. First, the *probability-integral transform* result, establishing that for any continuous variable X with distribution F_X , $F_X(X) \sim \mathcal{U}_{[0,1]}$, and second, *Sklar’s theorem* (Sklar, 1959), stating that for any continuous bivariate distribution whose cdf is F_{XY} , there exists a unique function C such that

$$F_{XY}(x, y) = C(F_X(x), F_Y(y)) \quad \forall (x, y) \in \mathbb{R}^2, \quad (1.1)$$

where F_X and F_Y are the marginals of F_{XY} . According to the previous definition, this function C is, indeed, a copula, called the copula of F_{XY} . From (1.1) it is clear that C describes how the two marginal distributions F_X and F_Y ‘interact’ to produce the joint F_{XY} . It, therefore, disjoints the marginal behaviours of X and Y from their dependence structure, hence the attractiveness of the copula approach. See Joe (1997) and Nelsen (2006) for book length treatment of the foregoing ideas. Other, more compact reviews include Genest and Favre (2007), Härdle and Okhrin (2009) and Embrechts (2009). Today, copulas are used extensively in statistical modelling in all areas, from quantitative finance and insurance to medicine and climatology. Therefore, empirically estimating a copula function from an observed bivariate sample $\{(X_i, Y_i); i = 1, \dots, n\}$ drawn from F_{XY} has become an important problem of modern statistical modelling.

Of course, estimating C essentially amounts to fitting a bivariate distribution, for what many parametric families have been suggested and studied: Gaussian, Student- t , Clayton, Frank or Gumbel copulas among others (see again Joe (1997) or Nelsen (2006) for details). These parametric models have formed the main body of the literature in the field so far. However, they suffer from the usual lack of flexibility of parametric approaches and the induced risk of misspecification. For instance, it has been argued that the main reason behind the 2009 global financial crisis was a reckless usage of the Gaussian copula (Salmon, 2009). There is, therefore, a tremendous need for flexible nonparametric copula models, making no rigid assumptions on the underlying distributions. An early step in that direction was the empirical copula devised by Deheuvels (1979). The related empirical copula process was studied further in Fermanian et

al (2004), Tsukahara (2005), Segers (2012) and Bücher and Volgushev (2013), and turns out to be the cornerstone of a variety of nonparametric copula-based procedures, see e.g. Genest and Rémillard (2004), Genest et al (2009a), Gudendorf and Segers (2012) or Li and Genton (2013), to cite only a few. Moreover, Fermanian and Scaillet (2003), Chen and Huang (2007), Omelka et al (2009) and Gijbels et al (2010) studied kernel methods to obtain flexible smooth estimates of the bivariate cdf C .

It is usually the case, though, that a distribution is more readily interpretable in terms of its probability density function than directly in terms of its cdf, and a copula is, in many aspects, no different. Assume that the bivariate cdf C is absolutely continuous. Then, its associated density is

$$c(u, v) = \frac{\partial^2 C}{\partial u \partial v}(u, v)$$

for $(u, v) \in \mathcal{I}$, a function naturally enough called the *copula density*. This paper precisely addresses the problem of estimating this copula density c in a nonparametric way, for what kernel methods again appear natural. This approach was pioneered in Behnen et al (1985) and Gijbels and Mielniczuk (1990), and arguably remains very attractive compared to its competitors, such as splines smoothing (Shen et al, 2008, Kauermann et al, 2013), wavelets (Hall and Neumeyer, 2006, Genest et al, 2009b, Autin et al, 2010), Bernstein polynomials (Bouezmarni et al, 2010, 2013, Janssen et al, 2013) or others (Qu and Yin, 2012), for its simplicity in construction, implementation and interpretation.

At least three factors make kernel estimation of $c(u, v)$ not standard, though, and have delayed the development of reliable kernel copula density estimators. A major concern is that kernel estimators suffer from boundary bias problems. Given the bivariate sample $\{(U_i = F_X(X_i), V_i = F_Y(Y_i)); i = 1, \dots, n\}$, the standard kernel estimator for c , say \hat{c}^* , at $(u, v) \in \mathcal{I}$ would be (Wand and Jones, 1995, Chapter 4)

$$\hat{c}^*(u, v) = \frac{1}{n|\mathbf{H}_{UV}|^{1/2}} \sum_{i=1}^n \mathbf{K} \left(\mathbf{H}_{UV}^{-1/2} \begin{pmatrix} u - U_i \\ v - V_i \end{pmatrix} \right), \quad (1.2)$$

where $\mathbf{K} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a bivariate kernel function and \mathbf{H}_{UV} is a symmetric positive-definite bandwidth matrix. It is, however, well known that an estimator such as (1.2) is in general not consistent on the boundaries of \mathcal{I} : it does not ‘feel’ the support boundaries of the underlying density and places through \mathbf{K} positive mass outside that support. In fact, standard kernel density arguments show that $\mathbb{E}(\hat{c}^*(u, v)) = \frac{1}{4}c(u, v) + O(h)$ at corners $((u, v) \in \{(0, 0), (1, 0), (0, 1), (1, 1)\})$ and $\mathbb{E}(\hat{c}^*(u, v)) = \frac{1}{2}c(u, v) + O(h)$ on the borders $((u, v) \in \{(0, \nu), (1, \nu), (\mu, 0), (\mu, 1) : \mu, \nu \in (0, 1)\})$. Although some papers ignored these boundary issues (Fermanian and Scaillet, 2003, Fermanian, 2005, Scaillet, 2007, Faugeras, 2009), it is clear that

accurate estimation of c calls for some boundary correction. Such corrections have indeed been proposed, inspired by ideas developed for univariate density estimation, e.g. mirror reflection (Gijbels and Mielniczuk, 1990) or the usage of boundary kernels (Chen and Huang, 2007), but with mixed results.

Secondly, kernel estimators are not consistent for unbounded densities. Yet, unlike most common probability densities, many copula densities of interest are unbounded. For instance, even in the apparently easy case of a bivariate Normal vector with moderate correlation, the copula density is unbounded in two of the corners of \mathcal{I} . It is, therefore, particularly important to use estimators able to cope with such unboundedness. Finally, estimating c cannot be made from a genuine random sample from its cdf C , as C is the distribution of $(U, V) = (F_X(X), F_Y(Y))$ and F_X and F_Y are typically unknown. Hence, the observations (U_i, V_i) are unavailable, and estimator (1.2) is, in fact, infeasible. In the copula literature, it is customary to use the ‘pseudo-observations’

$$\hat{U}_i = \frac{n}{n+1} \hat{F}_{X_n}(X_i) \quad \text{and} \quad \hat{V}_i = \frac{n}{n+1} \hat{F}_{Y_n}(Y_i) \quad (1.3)$$

where $\hat{F}_{X_n}(x) = \frac{1}{n} \sum_{j=1}^n \mathbb{I}_{\{X_j \leq x\}}$ is the empirical cdf of X , and similarly for \hat{F}_{Y_n} . The rescaling by $n/(n+1)$ in (1.3), aiming at keeping the pseudo-observations in the interior of $[0, 1]$, is also common practice. Then, the pseudo-sample $\{(\hat{U}_i, \hat{V}_i); i = 1, \dots, n\}$ is treated mostly as a sample from C and used instead of the ‘true’ sample $\{(U_i, V_i); i = 1, \dots, n\}$, although this may affect the statistical properties of the ensuing estimators (Charpentier et al, 2007, Genest and Segers, 2010).

The aim of this paper is to propose and study a new, kernel-type estimator of the copula density c . It is, in fact, the extension to the copula density case of the kernel-type estimator for univariate densities supported on the unit interval recently suggested in Geenens (2014). That estimator takes the constrained nature of the support into account from the outset, i.e. without relying on *ad hoc* boundary corrections (reflection, boundary kernels, etc.). It proved superior to its main competitors in the simulation studies for a wide range of density shapes, including for unbounded densities. The idea seems, therefore, suitable for estimating copula densities as well. Specifically, Geenens (2014)’s estimator makes use of the transformation method, building on ideas first suggested in Devroye and Györfi (1985, Chapter 9) and Marron and Ruppert (1994). In short, the initial $[0, 1]$ -supported variable of interest is transformed through the probit function into a variable whose support is unconstrained, the density of that transformed variable is estimated and an estimate of the initial density on $[0, 1]$ is obtained by back-transformation. This method appears very natural and yields very good results, provided that the estimation step in the transformed domain is carried out with care, as Geenens (2014)’s results showed.

Exploring this idea in more details in the context of copula density estimation is the topic of Section 2. Several versions of the estimators will be suggested, and their asymptotic properties will be derived in Section 3. Section 4 will address the crucial point of smoothing parameter selection in this framework. Simulation studies evidencing the very good practical behaviour of the probit-transformation estimators (Section 5), a real data analysis (Section 6) and some final remarks (Section 7) conclude the paper.

2 Probit transformation kernel copula density estimation

2.1 Probit transformation

As recalled above, direct kernel estimation of the density c of (U, V) is made difficult mainly by the constrained nature of its support $\mathcal{I} = [0, 1]^2$. Now, define

$$S = \Phi^{-1}(U) \quad \text{and} \quad T = \Phi^{-1}(V),$$

where Φ is the standard normal cdf and Φ^{-1} is its quantile function (i.e. the probit transformation). Given that both U and V are $\mathcal{U}_{[0,1]}$, S and T both follow standard normal distributions, which, nevertheless, does not imply that the vector (S, T) is bivariate normal. That will only be the case if the copula of the joint cdf of (S, T) , say F_{ST} , is the Gaussian copula, that is, if the copula C of F_{XY} itself is the Gaussian copula, as copulas are invariant to increasing transformations of their margins (Nelsen, 2006, Theorem 2.4.3). The idea is that, if $c(u, v) > 0$ Lebesgue-a.e. over \mathcal{I} (which will be assumed throughout the paper), (S, T) has unconstrained support \mathbb{R}^2 and estimating its density f_{ST} cannot suffer from boundary issues. In addition, due to its normal margins, one can expect f_{ST} to be well-behaved, and its estimation easy. In particular, under mild assumptions, f_{ST} and its partial derivatives up to the second order will be seen to be uniformly bounded on \mathbb{R}^2 , even in the case of unbounded copula density c (Lemma A.1 in the Appendix).

As the copula of F_{ST} is C , $S \sim \mathcal{N}(0, 1)$ and $T \sim \mathcal{N}(0, 1)$, one can write Sklar's theorem (1.1) for (S, T) :

$$F_{ST}(s, t) = \mathbb{P}(S \leq s, T \leq t) = C(\Phi(s), \Phi(t)), \quad \forall (s, t) \in \mathbb{R}^2.$$

Upon differentiation with respect to s and t , the joint density f_{ST} of (S, T) is found to be

$$f_{ST}(s, t) = c(\Phi(s), \Phi(t))\phi(s)\phi(t), \tag{2.1}$$

where ϕ is the standard normal density. Inverting this expression, one obtains

$$c(u, v) = \frac{f_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} \quad (2.2)$$

for any $(u, v) \in (0, 1)^2$. So, any estimator \hat{f}_{ST} of f_{ST} on \mathbb{R}^2 automatically produces an estimator of the copula density on the interior of \mathcal{I} , viz.

$$\hat{c}^{(\tau)}(u, v) = \frac{\hat{f}_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} \quad (2.3)$$

where the superscript (τ) refers to the idea of transformation. When necessary, $\hat{c}^{(\tau)}$ can also be defined at the boundaries of \mathcal{I} by continuity. This estimator enjoys interesting properties. Clearly, $\hat{c}^{(\tau)}$ cannot allocate any positive probability weight outside \mathcal{I} , since $(\Phi^{-1}(u), \Phi^{-1}(v))$ is not defined for $(u, v) \notin \mathcal{I}$. Also, if \hat{f}_{ST} is a *bona fide* density function, in the sense that $\hat{f}_{ST}(s, t) \geq 0$ for all (s, t) and $\iint_{\mathbb{R}^2} \hat{f}_{ST}(s, t) ds dt = 1$, then automatically $\hat{c}^{(\tau)}(u, v) \geq 0$ for all $(u, v) \in \mathcal{I}$ and $\iint_{\mathcal{I}} \hat{c}^{(\tau)}(u, v) du dv = 1$. This is easily seen through the changes of variable $u = \Phi(s)$ and $v = \Phi(t)$. Finally, if \hat{f}_{ST} is a uniformly (weak or strong) consistent estimator for f_{ST} , i.e. $\sup_{(s,t) \in \mathbb{R}^2} |\hat{f}_{ST}(s, t) - f_{ST}(s, t)| \rightarrow 0$ in probability or almost surely as $n \rightarrow \infty$, the estimator $\hat{c}^{(\tau)}$ inherits that uniform consistency on any compact proper subset of \mathcal{I} .

2.2 The naive estimator

A first natural idea would be to use the standard kernel density estimator as \hat{f}_{ST} in (2.3). Specifically, one would like to use an estimator like

$$\hat{f}_{ST}^*(s, t) = \frac{1}{n|\mathbf{H}_{ST}|^{1/2}} \sum_{i=1}^n \mathbf{K} \left(\mathbf{H}_{ST}^{-1/2} \begin{pmatrix} s - S_i \\ t - T_i \end{pmatrix} \right) \quad (2.4)$$

where \mathbf{K} is a bivariate kernel function and \mathbf{H}_{ST} is some symmetric positive-definite bandwidth matrix, and $\{(S_i = \Phi^{-1}(U_i), T_i = \Phi^{-1}(V_i)); i = 1, \dots, n\}$ is the sample in the transformed domain. However, as the (U_i, V_i) 's are unavailable in this context, so are the (S_i, T_i) 's. Instead, one has to use

$$\{(\hat{S}_i = \Phi^{-1}(\hat{U}_i), \hat{T}_i = \Phi^{-1}(\hat{V}_i)); i = 1, \dots, n\}, \quad (2.5)$$

the pseudo-transformed sample. The feasible version of (2.4) is, therefore,

$$\hat{f}_{ST}(s, t) = \frac{1}{n|\mathbf{H}_{ST}|^{1/2}} \sum_{i=1}^n \mathbf{K} \left(\mathbf{H}_{ST}^{-1/2} \begin{pmatrix} s - \hat{S}_i \\ t - \hat{T}_i \end{pmatrix} \right). \quad (2.6)$$

Through (2.3), this directly leads to the following probit transformation kernel copula density estimator:

$$\hat{c}^{(\tau)}(u, v) = \frac{1}{n|\mathbf{H}_{ST}|^{1/2}\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} \sum_{i=1}^n \mathbf{K} \left(\mathbf{H}_{ST}^{-1/2} \begin{pmatrix} \Phi^{-1}(u) - \Phi^{-1}(\hat{U}_i) \\ \Phi^{-1}(v) - \Phi^{-1}(\hat{V}_i) \end{pmatrix} \right). \quad (2.7)$$

This is essentially the estimator suggested in Charpentier et al (2007), also used as-is in Lopez-Paz et al (2013), although it was not studied in any details in those two papers. Omelka et al (2009) derived the theoretical properties of an estimator for the copula C (not its density) based on the same transformation.

This idea was, however, called ‘naive’ in Geenens (2014) in the univariate case. In fact, the method does not provide good results close to the boundaries, even though it was designed to fix boundary issues. Indeed, the estimator (2.7) will be seen not to perform well in the next sections. Geenens (2014) explained the reasons for that failure, and suggested some remedies. In particular, estimating the density in the transformed domain via local likelihood methods (Loader, 1996, Hjort and Jones, 1996) offers a promising alternative while keeping the simplicity and the intuitive appeal of the probit-transformation estimator. This is investigated for estimating a copula density in the next subsection.

2.3 Improved probit-transformation copula density estimators

Loader (1996) and Hjort and Jones (1996) proposed two different, although similar in many aspects, formulations of the local likelihood density estimator. Loader (1996) locally approximates the logarithm of the unknown density by a polynomial, whereas Hjort and Jones (1996) consider local parametric density modelling. This paper will only make use of Loader (1996)’s idea, mainly because the asymptotic theory is more transparent. In any case, both formulations share most of their advantages and drawbacks, and typically yield very similar estimates.

In this setting of estimating f_{ST} from the pseudo-sample $\{(\hat{S}_i, \hat{T}_i); i = 1, \dots, n\}$, Loader (1996)’s local likelihood estimator is defined as follows. Around $(s, t) \in \mathbb{R}^2$, $\log f_{ST}$ is assumed to be well approximated by a polynomial of some order p . Classically, only local log-linear ($p = 1$) and local log-quadratic ($p = 2$) estimators are considered. Specifically, in the first case ($p = 1$), it is assumed that, for (\check{s}, \check{t}) ‘close’ to (s, t) ,

$$\log f_{ST}(\check{s}, \check{t}) \simeq a_{1,0}(s, t) + a_{1,1}(s, t)(\check{s} - s) + a_{1,2}(s, t)(\check{t} - t) \doteq P_{\mathbf{a}_1}(\check{s} - s, \check{t} - t) \quad (2.8)$$

and in the second case ($p = 2$)

$$\begin{aligned} \log f_{ST}(\check{s}, \check{t}) &\simeq a_{2,0}(s, t) + a_{2,1}(s, t)(\check{s} - s) + a_{2,2}(s, t)(\check{t} - t) \\ &\quad + a_{2,3}(s, t)(\check{s} - s)^2 + a_{2,4}(s, t)(\check{t} - t)^2 + a_{2,5}(s, t)(\check{s} - s)(\check{t} - t) \\ &\doteq P_{\mathbf{a}_2}(\check{s} - s, \check{t} - t). \end{aligned}$$

The vectors $\mathbf{a}_1(s, t) = (a_{1,0}(s, t), a_{1,1}(s, t), a_{1,2}(s, t))$ and $\mathbf{a}_2(s, t) \doteq (a_{2,0}(s, t), \dots, a_{2,5}(s, t))$ are then estimated by solving a weighted maximum likelihood problem. For either $p = 1, 2$,

$$\begin{aligned} \tilde{\mathbf{a}}_p(s, t) = \arg \max_{\mathbf{a}_p} \left\{ \sum_{i=1}^n \mathbf{K} \left(\mathbf{H}_{ST}^{-1/2} \begin{pmatrix} s - \hat{S}_i \\ t - \hat{T}_i \end{pmatrix} \right) P_{\mathbf{a}_p}(\hat{S}_i - s, \hat{T}_i - t) \right. \\ \left. - n \iint_{\mathbb{R}^2} \mathbf{K} \left(\mathbf{H}_{ST}^{-1/2} \begin{pmatrix} s - \check{s} \\ t - \check{t} \end{pmatrix} \right) \exp(P_{\mathbf{a}_p}(\check{s} - s, \check{t} - t)) d\check{s} d\check{t} \right\}, \quad (2.9) \end{aligned}$$

where, as previously, \mathbf{K} is a bivariate kernel function and \mathbf{H}_{ST} is a symmetric positive-definite bandwidth matrix. The estimate of f_{ST} at (s, t) is then, naturally, $\tilde{f}_{ST}^{(1)}(s, t) = \exp(\tilde{a}_{1,0}(s, t))$ for local log-linear, and $\tilde{f}_{ST}^{(2)}(s, t) = \exp(\tilde{a}_{2,0}(s, t))$ for local log-quadratic modelling. ‘Improved’ probit-transformation kernel copula density estimators for $c(u, v)$ are finally obtained through (2.3) as

$$\tilde{c}^{(\tau, p)}(u, v) = \frac{\tilde{f}_{ST}^{(p)}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} \quad (2.10)$$

for $p = 1$ and $p = 2$. The motivation and the advantages of estimating f_{ST} by local likelihood methods instead of raw kernel density estimation are related to the detailed discussion in Geenens (2014), and are therefore omitted here. The asymptotic properties of these estimators (‘naive’ and ‘improved’ probit-transformation kernel copula density estimators) are derived in the next section.

3 Asymptotic properties

For simplicity, it will be assumed that \mathbf{K} is a product Gaussian kernel, i.e. $\mathbf{K}(z_1, z_2) = \phi(z_1)\phi(z_2)$, and $\mathbf{H}_{ST} = h^2\mathbf{I}$ for some $h > 0$. Note that, in practice, there are reasons to keep an unconstrained, non-diagonal bandwidth matrix \mathbf{H}_{ST} . In particular, the copula density is typically stretched along one of the diagonals of the unit square when some dependence is present in (X, Y) , which provides a density f_{ST} likewise stretched along one of the 45 degrees lines in \mathbb{R}^2 . Hence, using a bandwidth matrix directing smoothing in that particular direction is sensible (Duong and Hazelton, 2005), as discussed further in Section 4. That said, theoretical results for that general case would be less tractable than, while qualitatively equivalent to,

the simpler case presented below. Note that for that particular kernel \mathbf{K} , $\iint \mathbf{K}^2(z_1, z_2) dz_1 dz_2 = (4\pi)^{-1}$ and $\iint z_k^2 \mathbf{K}(z_1, z_2) dz_1 dz_2 = 1$, $k = 1, 2$. These quantities frequently arise in the properties of kernel estimators, and direct use of these particular numerical values will be made in the results below.

3.1 The naive estimator and an amended version

Consider the naive estimator (2.7) which, with the above specifications of \mathbf{K} and \mathbf{H}_{ST} , reduces to

$$\hat{c}^{(\tau)}(u, v) = \frac{1}{nh^2 \phi(\Phi^{-1}(u)) \phi(\Phi^{-1}(v))} \sum_{i=1}^n \phi\left(\frac{\Phi^{-1}(u) - \Phi^{-1}(\hat{U}_i)}{h}\right) \phi\left(\frac{\Phi^{-1}(v) - \Phi^{-1}(\hat{V}_i)}{h}\right). \quad (3.1)$$

Given (2.3), it is clear that its statistical properties will entirely depend on those of (2.6), here

$$\hat{f}_{ST}(s, t) = \frac{1}{nh^2} \sum_{i=1}^n \phi\left(\frac{s - \hat{S}_i}{h}\right) \phi\left(\frac{t - \hat{T}_i}{h}\right). \quad (3.2)$$

If f_{ST} admits continuous second-order partial derivatives, expressions for the bias and the variance of the ideal, infeasible estimator \hat{f}_{ST}^* (2.4), as well as its asymptotic normality, are well known (Wand and Jones, 1995, Chapter 4). Proposition 3.1 below ascertains that using the pseudo-observations (2.5) instead of genuine ones does not affect those properties. Note that (3.2) can be written

$$\hat{f}_{ST}(s, t) = \frac{1}{h^2} \iint_{\mathbb{R}^2} \phi\left(\frac{s - \Phi^{-1}(u)}{h}\right) \phi\left(\frac{t - \Phi^{-1}(v)}{h}\right) d\hat{C}_n(u, v), \quad (3.3)$$

where \hat{C}_n is the empirical copula

$$\hat{C}_n(u, v) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{U}_i \leq u, \hat{V}_i \leq v\}}. \quad (3.4)$$

Hence, although living in the transformed domain, the behaviour of $\hat{f}_{ST}(s, t)$ will be driven by the properties of \hat{C}_n on \mathcal{I} . The following assumptions will be made.

Assumption 3.1. *The sample $\{(X_i, Y_i); i = 1, \dots, n\}$ is an i.i.d. sample from the joint distribution F_{XY} , an absolutely continuous distribution with marginals F_X and F_Y strictly increasing on their support;*

Assumption 3.2. *The copula C of F_{XY} is such that $(\partial C / \partial u)(u, v)$ and $(\partial^2 C / \partial u^2)(u, v)$ exist and are continuous on $\{(u, v) : u \in (0, 1), v \in [0, 1]\}$, and $(\partial C / \partial v)(u, v)$ and $(\partial^2 C / \partial v^2)(u, v)$ exist and are continuous on $\{(u, v) : u \in [0, 1], v \in (0, 1)\}$. In addition, there are constants K_1 and K_2 such that*

$$\left\{ \begin{array}{ll} \left| \frac{\partial^2 C}{\partial u^2}(u, v) \right| \leq \frac{K_1}{u(1-u)} & \text{for } (u, v) \in (0, 1) \times [0, 1]; \\ \left| \frac{\partial^2 C}{\partial v^2}(u, v) \right| \leq \frac{K_2}{v(1-v)} & \text{for } (u, v) \in [0, 1] \times (0, 1); \end{array} \right.$$

Assumption 3.3. *The density c of C exists, is positive and admits continuous second-order partial derivatives on the interior of the unit square \mathcal{I} . In addition, there is a constant K_{00} such that*

$$c(u, v) \leq K_{00} \min \left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)} \right) \quad \forall (u, v) \in (0, 1)^2. \quad (3.5)$$

Assumption 3.1 guarantees the existence and the uniqueness of the copula C of F_{XY} . Assumptions 3.2-3.3 mostly reduce to Conditions 2.1 and 4.1 in Segers (2012), who claims that they hold for many copula families, such as Gaussian, Archimedean and most extreme-value copulas. Moreover, Omelka et al (2009) explicitly show that they are satisfied by the Clayton, Gumbel, Gaussian and Student copulas. Compared to Segers (2012), Assumption 3.3 only requires further the existence and continuity of second-order partial derivatives of c , which is natural in kernel estimation. It is worth noting that c is allowed to grow unboundedly in some of the corners of \mathcal{I} , provided (3.5) remains valid. The following result can now be stated. An important observation is that it holds true for $h \sim n^{-a}$, $a \in (0, 1/4)$, which includes the optimal bandwidth order known to be $h \sim n^{-1/6}$ for bivariate density estimation.

Proposition 3.1. *Assume that $\mathbf{K}(z_1, z_2) = \phi(z_1)\phi(z_2)$ and $\mathbf{H}_{ST} = h^2\mathbf{I}$ with $h \sim n^{-a}$ for some $a \in (0, 1/4)$.*

Under Assumptions 3.1-3.3, the estimator (3.2) at any $(s, t) \in \mathbb{R}^2$ is such that

$$\sqrt{nh^2} \left(\hat{f}_{ST}(s, t) - f_{ST}(s, t) - h^2 b_{ST}(s, t) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma_{ST}^2(s, t) \right), \quad (3.6)$$

where $b_{ST}(s, t) = \frac{1}{2} \left(\frac{\partial^2 f_{ST}}{\partial s^2}(s, t) + \frac{\partial^2 f_{ST}}{\partial t^2}(s, t) \right)$ and $\sigma_{ST}^2(s, t) = \frac{f_{ST}(s, t)}{4\pi}$.

Proof. See Appendix. □

As recalled in Section 1, resorting to pseudo-observations is known to usually affect the statistical properties of the estimators of interest in copula modelling. In particular, an overriding result in the field is the weak convergence of the empirical copula process

$$\mathbb{C}_n(u, v) \doteq \sqrt{n}(\hat{C}_n(u, v) - C(u, v)) \rightsquigarrow \mathbb{G}_C(u, v) \doteq \mathbb{B}_C(u, v) - \frac{\partial C}{\partial u}(u, v)\mathbb{B}_C(u, 1) - \frac{\partial C}{\partial v}(u, v)\mathbb{B}_C(1, v), \quad (3.7)$$

where $\mathbb{B}_C(u, v)$ is a bivariate pinned C -Brownian Sheet on \mathcal{I} , i.e. the tight centred Gaussian process whose covariance function is $\mathbb{E}(\mathbb{B}_C(u, v)\mathbb{B}_C(u', v')) = C(u \wedge u', v \wedge v') - C(u, v)C(u', v')$ (Fermanian et al, 2004, Segers, 2012). In fact, $\mathbb{B}_C(u, v)$ would be the limiting process if the margins were known, i.e. if ‘genuine’ U_i ’s and V_i ’s were used in (3.4). The extra two terms in the right-hand side of (3.7) are, therefore, often interpreted as ‘the price to pay’ for using pseudo-observations – although this effect is sometimes advantageous (Genest and Segers, 2010). Yet, the proof of Proposition 3.1 reveals that the effect of those

two terms asymptotically vanishes when one looks at the properties of the kernel density estimator (3.2). As a result, the rate of convergence, as well as the expressions for asymptotic bias and variance, are the same as what one would obtain for the ideal estimator \hat{f}_{ST}^* using genuine i.i.d. observations from C . Intuitively, this is because a kernel density estimator converges slower than an empirical distribution function. Resorting to pseudo-observations may disturb the \sqrt{n} -convergence of the latter, but it goes unnoticed (asymptotically) compared to the nonparametric convergence rate $O((nh^2)^{-1/2})$ of the former.

Now, differentiating (2.1) yields

$$\frac{\partial f_{ST}}{\partial s}(s, t) = \frac{\partial c}{\partial u}(\Phi(s), \Phi(t))\phi^2(s)\phi(t) - sc(\Phi(s), \Phi(t))\phi(s)\phi(t), \quad \text{and} \quad (3.8)$$

$$\frac{\partial^2 f_{ST}}{\partial s^2}(s, t) = \frac{\partial^2 c}{\partial u^2}(\Phi(s), \Phi(t))\phi^3(s)\phi(t) - 3s \frac{\partial c}{\partial u}(\Phi(s), \Phi(t))\phi^2(s)\phi(t) + (s^2 - 1)c(\Phi(s), \Phi(t))\phi(s)\phi(t) \quad (3.9)$$

(and similar for $\frac{\partial f_{ST}}{\partial t}$, $\frac{\partial^2 f_{ST}}{\partial t^2}$ and $\frac{\partial^2 f_{ST}}{\partial s \partial t}$). Hence, combining (2.1), (2.3), (3.6) and (3.9), one can state the following theorem for the ‘naive’ probit transformation kernel copula density estimator (3.1).

Theorem 3.1. *Under the assumptions of Proposition 3.1, the ‘naive’ probit transformation kernel copula density estimator (3.1) at any $(u, v) \in (0, 1)^2$ is such that*

$$\sqrt{nh^2} \left(\hat{c}^{(\tau)}(u, v) - c(u, v) - h^2 b(u, v) \right) \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma^2(u, v)),$$

where

$$b(u, v) = \frac{1}{2} \left\{ \frac{\partial^2 c}{\partial u^2}(u, v)\phi^2(\Phi^{-1}(u)) + \frac{\partial^2 c}{\partial v^2}(u, v)\phi^2(\Phi^{-1}(v)) \right. \\ \left. - 3 \left(\frac{\partial c}{\partial u}(u, v)\Phi^{-1}(u)\phi(\Phi^{-1}(u)) + \frac{\partial c}{\partial v}(u, v)\Phi^{-1}(v)\phi(\Phi^{-1}(v)) \right) \right. \\ \left. + c(u, v) (\{\Phi^{-1}(u)\}^2 + \{\Phi^{-1}(v)\}^2 - 2) \right\} \quad (3.10)$$

$$\text{and } \sigma^2(u, v) = \frac{c(u, v)}{4\pi\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))}.$$

Proof. See Appendix. □

It is seen that, when (u, v) approaches one of the boundaries, both the (asymptotic) bias and variance of the estimator tend to grow unboundedly. Indeed, $\sigma^2(u, v) \propto c(u, v)/(\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v)))$ and $b(u, v)$ includes the term $c(u, v) (\{\Phi^{-1}(u)\}^2 + \{\Phi^{-1}(v)\}^2 - 2)$, and the functions $\Phi^{-1}(\cdot)$ and $1/\phi(\Phi^{-1}(\cdot))$ are unbounded. Thus, along the boundaries, $\hat{c}^{(\tau)}$ will work properly only over areas, if any, where c approaches 0 very




Figure 3.1: True Gaussian copula density with $\rho = 0.3$ (left), its naive probit-transformation kernel estimator from a typical random sample of size $n = 1000$ (middle) and its amended naive probit-transformation kernel estimator from the same sample (right). The (unconstrained) bandwidth matrix \mathbf{H}_{ST} was chosen by the Normal Reference rule in the (S, T) -domain.

smoothly. Otherwise, $\hat{c}^{(\tau)}$ will typically show a very erratic behaviour (large variance) and will be prone to exploding (large positive bias), especially in the corners of \mathcal{I} in which c is large. Figure 3.1 illustrates these problems, from a typical sample of size $n = 1000$ simulated from the Gaussian copula with correlation $\rho = 0.3$ (left panel). The corresponding naive probit-transformation kernel estimator is shown in the middle panel. An unconstrained matrix \mathbf{H}_{ST} was used in (2.6)/(2.7) and chosen by the multivariate Normal Reference rule (Chacón et al, 2011). Here, this is optimal: C being a Gaussian copula, f_{ST} is a bivariate normal density. Over the middle of the unit square, the estimator appears to work decently, but towards the boundaries the estimate shows coarse folds and, indeed, hypertrophies the peaks in the corners $(0, 0)$ and $(1, 1)$. Clearly, this estimator is not acceptable as-is. It is, therefore, not surprising that it has been reported not to perform well, see for instance the simulation study in Bouezmarni et al (2013).

The third, unbounded term in (3.10) can, however, be easily adjusted for. Instead of (2.3), take

$$\hat{c}^{(\tau\text{am})}(u, v) = \frac{\hat{f}_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} \times \frac{1}{1 + \frac{1}{2}h^2 (\{\Phi^{-1}(u)\}^2 + \{\Phi^{-1}(v)\}^2 - 2)}. \quad (3.11)$$

For this ‘amended’ version of $\hat{c}^{(\tau)}$, one can see that the asymptotic bias becomes proportional to

$$b^{(\text{am})}(u, v) = \frac{1}{2} \left\{ \frac{\partial^2 c}{\partial u^2}(u, v)\phi^2(\Phi^{-1}(u)) + \frac{\partial^2 c}{\partial v^2}(u, v)\phi^2(\Phi^{-1}(v)) - 3 \left(\frac{\partial c}{\partial u}(u, v)\Phi^{-1}(u)\phi(\Phi^{-1}(u)) + \frac{\partial c}{\partial v}(u, v)\Phi^{-1}(v)\phi(\Phi^{-1}(v)) \right) \right\}.$$

In fact, the deterministic, multiplicative amendment in (3.11) exactly makes it up for the third term in (3.10) in the asymptotic development, given that $(1 + h^2)^{-1} = 1 - h^2 + o(h^2)$ as $h \rightarrow 0$. The improvement is

illustrated in Figure 3.1 (right panel), where the amended version of the naive estimator computed on the same data set as in the middle panel is shown. The peaks in the corners $(0, 0)$ and $(1, 1)$ are now roughly of the right height. On the other hand, the wiggly appearance of the estimate along boundaries mostly remains, as the variance is not affected by the deterministic amendment. On a side note, the amendment implies that the estimator $\hat{c}^{(\tau\text{am})}(u, v)$ does not integrate to one over the unit square any more, which calls for a renormalisation such as $\hat{c}^{(\tau\text{am})}(u, v) \leftarrow \hat{c}^{(\tau\text{am})}(u, v) / \iint_{\mathcal{I}} \hat{c}^{(\tau\text{am})}(u, v) du dv$. This is, however, frequent in other nonparametric density estimation procedures, and is not really a problem.

3.2 Improved probit-transformation kernel copula density estimators

Now the asymptotic properties of the ‘improved’ versions of the probit-transformation kernel copula density estimators are derived. Again, for convenience, the results are stated in the case where \mathbf{K} is a product of two univariate Gaussian kernels and $\mathbf{H}_{ST} = h^2 \mathbf{I}$, for some $h > 0$, in (2.9). The first version estimates the joint density f_{ST} in the transformed domain by the local log-linear estimator $\tilde{f}_{ST}^{(1)}$. Consider first the ‘ideal’ version $\tilde{f}_{ST}^{*(1)}$ of this estimator, computed on the genuine sample $\{(S_i, T_i), i = 1, \dots, n\}$. From Loader (1996), one gets, for all $(s, t) \in \mathbb{R}^2$ at which $f_{ST}(s, t)$ is positive and admits continuous second-order partial derivatives,

$$\sqrt{nh^2} \left(\tilde{f}_{ST}^{*(1)}(s, t) - f_{ST}(s, t) - h^2 b_{ST}^{(1)}(s, t) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma_{ST}^{(1)2}(s, t) \right), \quad (3.12)$$

where

$$b_{ST}^{(1)}(s, t) = \frac{1}{2} \left\{ \left(\frac{\partial^2 f_{ST}}{\partial s^2} + \frac{\partial^2 f_{ST}}{\partial t^2} \right) (s, t) - \frac{1}{f_{ST}(s, t)} \left(\left\{ \frac{\partial f_{ST}}{\partial s} \right\}^2 + \left\{ \frac{\partial f_{ST}}{\partial t} \right\}^2 \right) (s, t) \right\}$$

and $\sigma_{ST}^{(1)2}(s, t) = \frac{f_{ST}(s, t)}{4\pi}$. Of course, if $f_{ST}(s, t) = 0$ at some (s, t) , the singularity of the log-density cannot be accurately approximated by (2.8), but this is ruled out here by Assumption 3.3 which requires c to be positive all over the interior of the unit square. By (2.1), this implies that f_{ST} is positive over \mathbb{R}^2 .

Define the ‘ideal’ local log-linear probit-transformation kernel copula density estimator $\tilde{c}^{*(\tau, 1)}(u, v) = \tilde{f}_{ST}^{*(1)}(\Phi^{-1}(u), \Phi^{-1}(v)) / (\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v)))$. By the same token as for Theorem 3.1, in particular by using (2.1), (2.3), (3.8) and (3.9) in (3.12), one can obtain

$$\sqrt{nh^2} \left(\tilde{c}^{*(\tau, 1)}(u, v) - c(u, v) - h^2 b^{(1)}(u, v) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma^{(1)2}(u, v) \right), \quad (3.13)$$

where

$$\begin{aligned}
b^{(1)}(u, v) = \frac{1}{2} & \left\{ \frac{\partial^2 c}{\partial u^2}(u, v) \phi^2(\Phi^{-1}(u)) + \frac{\partial^2 c}{\partial v^2}(u, v) \phi^2(\Phi^{-1}(v)) \right. \\
& - \frac{1}{c(u, v)} \left(\left\{ \frac{\partial c}{\partial u}(u, v) \right\}^2 \phi^2(\Phi^{-1}(u)) + \left\{ \frac{\partial c}{\partial v}(u, v) \right\}^2 \phi^2(\Phi^{-1}(v)) \right) \\
& \left. - \left(\frac{\partial c}{\partial u}(u, v) \Phi^{-1}(u) \phi(\Phi^{-1}(u)) + \frac{\partial c}{\partial v}(u, v) \Phi^{-1}(v) \phi(\Phi^{-1}(v)) \right) - 2c(u, v) \right\} \quad (3.14)
\end{aligned}$$

and $\sigma^{(1)2}(u, v) = \frac{c(u, v)}{4\pi\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))}$.

The next result ascertains that, like for the ‘naive’ estimator, the asymptotic properties of $\tilde{c}^{(\tau,1)}$ are not affected by using pseudo-observations, and are consequently identical to those of the ideal version $\tilde{c}^{*(\tau,1)}$.

Theorem 3.2. *Under the assumptions of Proposition 3.1, the ‘improved’ local log-linear probit-transformation kernel copula density estimator $\tilde{c}^{(\tau,1)}$ at any $(u, v) \in (0, 1)^2$ is such that*

$$\sqrt{nh^2} \left(\tilde{c}^{(\tau,1)}(u, v) - c(u, v) - h^2 b^{(1)}(u, v) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma^{(1)2}(u, v) \right),$$

where $b^{(1)}(u, v)$ and $\sigma^{(1)2}(u, v)$ are given above.

Proof. See Appendix. □

Compared to the ‘naive’ estimator, the variance is the same but the bias is significantly different. It is now automatically free from any unbounded terms. In fact, Hjort and Jones (1996) showed (their expression (7.3)) that the local log-linear and the standard kernel estimators in the (S, T) -domain ($\tilde{f}_{ST}^{(1)}$ and \hat{f}_{ST} , respectively) satisfy

$$\tilde{f}_{ST}^{(1)}(s, t) = \hat{f}_{ST}(s, t) \exp \left\{ -\frac{1}{2} h^2 \left[\left(\frac{\partial \hat{f}_{ST}(s, t) / \partial s}{\hat{f}_{ST}(s, t)} \right)^2 + \left(\frac{\partial \hat{f}_{ST}(s, t) / \partial t}{\hat{f}_{ST}(s, t)} \right)^2 \right] \right\}. \quad (3.15)$$

This closed-form expression for $\tilde{f}_{ST}^{(1)}$ shows that it improves on the basic kernel estimator by adjusting for the local slopes. With (2.3), (2.10) and an analogue of (3.8) for hat versions, one can state a similar result




Figure 3.2: True Gaussian copula density with $\rho = 0.3$ (left), its local log-linear (middle) and log-quadratic (right) improved probit-transformation kernel estimators from the same sample ($n = 500$) as in Figure 3.1. Both estimates use an unconstrained bandwidth matrix \mathbf{H}_{ST} chosen by cross-validation in the (S, T) -domain, see Section 4.

in terms of the copula density estimators:

$$\begin{aligned} \tilde{c}^{(\tau,1)}(u, v) &= \hat{c}(u, v) \\ &\times \exp \left\{ -\frac{1}{2}h^2 \left[\left(\frac{\partial \hat{c}(u, v)/\partial u}{\hat{c}(u, v)} \right)^2 \phi^2(\Phi^{-1}(u)) + \left(\frac{\partial \hat{c}(u, v)/\partial v}{\hat{c}(u, v)} \right)^2 \phi^2(\Phi^{-1}(v)) \right. \right. \\ &\quad \left. \left. - 2 \left\{ \left(\frac{\partial \hat{c}(u, v)/\partial u}{\hat{c}(u, v)} \right) \Phi^{-1}(u) \phi(\Phi^{-1}(u)) + \left(\frac{\partial \hat{c}(u, v)/\partial v}{\hat{c}(u, v)} \right) \Phi^{-1}(v) \phi(\Phi^{-1}(v)) \right\} \right. \right. \\ &\quad \left. \left. + \{ \Phi^{-1}(u) \}^2 + \{ \Phi^{-1}(v) \}^2 \right] \right\}. \end{aligned}$$

This reveals that, not only the local log-linear estimator $\tilde{c}^{(\tau,1)}$ attempts a correction for the slopes of c like in (3.15), it actively acts on the boundary behaviour as well. Indeed, given that $\phi^2(\Phi^{-1}(\cdot))$ and $\Phi^{-1}(\cdot)\phi(\Phi^{-1}(\cdot))$ tend to 0 towards 0 and 1, the first four terms in the bracket in the previous expression will have little influence towards the boundaries (provided c does not tend to 0 too sharply there). On the other hand, $\{ \Phi^{-1}(u) \}^2 + \{ \Phi^{-1}(v) \}^2$ tends to $+\infty$ very fast along boundaries (and all the more in the corners), hence $\hat{c}(u, v)$ is multiplied by something quickly tending to 0 there and this prevents it from exploding. This is, in fact, very similar to what the amendment in (3.11) attempted, but is now achieved automatically. Figure 3.2 (middle panel) shows the estimate $\tilde{c}^{(\tau,1)}$ for the data set used in Figure 3.1. It used the cross-validation criterion discussed in Section 4 to select the matrix \mathbf{H}_{ST} in (2.9).

The second improved probit-transformation estimator is obtained when taking $p = 2$ in (2.9). Again, consider first the ‘ideal’ estimator $\tilde{f}_{ST}^{*(2)}$, computed on the genuine sample $\{(S_i, T_i); i = 1, \dots, n\}$. Locally fitting a polynomial of a higher degree is known to reduce the asymptotic bias of the estimator, here from order $O(h^2)$ to order $O(h^4)$ (Loader, 1996, Hjort and Jones, 1996), sufficient smoothness of f_{ST} permitting.

Specifically, if f_{ST} admits continuous fourth-order partial derivatives and is positive at (s, t) , then

$$\sqrt{nh^2} \left(\tilde{f}_{ST}^{*(2)}(s, t) - f_{ST}(s, t) - h^4 b_{ST}^{(2)}(s, t) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma_{ST}^{(2)2}(s, t) \right), \quad (3.16)$$

where $\sigma_{ST}^{(2)2}(s, t) = \frac{5}{2} \frac{f_{ST}(s, t)}{4\pi}$ and

$$b_{ST}^{(2)}(s, t) = -\frac{1}{8} f_{ST}(s, t) \times \left\{ \left(\frac{\partial^4 g}{\partial s^4} + \frac{\partial^4 g}{\partial t^4} \right) + 4 \left(\frac{\partial^3 g}{\partial s^3} \frac{\partial g}{\partial s} + \frac{\partial^3 g}{\partial t^3} \frac{\partial g}{\partial t} + \frac{\partial^3 g}{\partial s^2 \partial t} \frac{\partial g}{\partial t} + \frac{\partial^3 g}{\partial s \partial t^2} \frac{\partial g}{\partial s} \right) + 2 \frac{\partial^4 g}{\partial s^2 \partial t^2} \right\} (s, t),$$

with $g(s, t) = \log f_{ST}(s, t)$. Starting from $g(s, t) = \log c(\Phi(s), \Phi(t)) + \log \phi(s) + \log \phi(t)$, tedious algebraic differentiation provides all partial derivatives of g up to order four in terms of c and its partial derivatives up to order four. Naturally, c will be assumed to admit continuous fourth-order partial derivatives.

Assumption 3.4. *The copula density $c(u, v) = (\partial^2 C / \partial u \partial v)(u, v)$ admits continuous fourth-order partial derivatives on the interior of the unit square \mathcal{I} .*

As previously, it readily follows from (3.16) that

$$\sqrt{nh^2} \left(\tilde{c}^{*(\tau, 2)}(u, v) - c(u, v) - h^4 b^{(2)}(u, v) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma^{(2)2}(u, v) \right),$$

where $\sigma^{(2)2}(u, v) = \frac{5}{2} \frac{c(u, v)}{4\pi \phi(\Phi^{-1}(u)) \phi(\Phi^{-1}(v))}$ and $b^{(2)}(u, v)$ is an expression of the same type as (3.14), this time involving the partial derivatives of c up to the fourth order. Like above, it can be shown that resorting to pseudo-observations does not affect these properties. This, however, requires a condition on the bandwidth ($h \sim n^{-a}$, $a \in (0, 1/6)$) slightly stronger than previously. Given that the bias order is reduced to $O(h^4)$, the optimal bandwidth order is now seen to be $h \sim n^{-1/10}$, so that the bandwidth requirement does still include that optimal order.

Theorem 3.3. *Under the assumptions of Proposition 3.1 and Assumption 3.4, if $h \sim n^{-a}$ with $a \in (0, 1/6)$ as $n \rightarrow \infty$, the ‘improved’ local log-quadratic probit-transformation kernel copula density estimator $\tilde{c}^{(\tau, 2)}$ at any $(u, v) \in (0, 1)^2$ is such that*

$$\sqrt{nh^2} \left(\tilde{c}^{(\tau, 2)}(u, v) - c(u, v) - h^4 b^{(2)}(u, v) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma^{(2)2}(u, v) \right),$$

where $b^{(2)}(u, v)$ and $\sigma^{(2)2}(u, v)$ are described above.

Proof. See Appendix. □

For seek of conciseness, the expression of $b^{(2)}(u, v)$ is not given here (it is made up of several dozens of terms). The interesting point about it, though, is that, unlike (3.14) which shows a last term $-c(u, v)$, all terms of $b^{(2)}(u, v)$ are proportional to $\{\Phi^{-1}(u)\}^\alpha\{\Phi^{-1}(v)\}^\beta\phi^\gamma(\Phi^{-1}(u))\phi^\delta(\Phi^{-1}(v))$, for some non-negative integral powers $\alpha, \beta, \gamma, \delta$, and all those functions tend to 0 as $u, v \rightarrow 0/1$. Hence, $b^{(2)}(u, v)$ may actually tend to 0 towards the boundaries, and the bias there be actually of order $o(h^4)$. Again, this will be the case where c does not tend to 0 too sharply when approaching the boundary. The expression of the variance is the same as that for $\hat{c}^{(\tau)}(u, v)$ and $\tilde{c}^{(\tau,1)}(u, v)$, except that it has been inflated by a factor $5/2$. This inflation factor is, also, a well-known feature in local polynomial modelling when fitting a higher-degree polynomial (Fan and Gijbels, 1996, Section 3.3.1).

Interestingly, *ad-hoc* techniques for reducing the bias of kernel estimators from $O(h^2)$ to $O(h^4)$, e.g. higher-order kernels or multiplicative adjustment, have long been an active research topic (Jones and Signorini, 1997). Yet, few of those methods have actually taken hold. The main reason is that the demonstrated improvement is an asymptotic result, which usually goes unnoticed for sample sizes one typically has in practice while implying interpretability issues (e.g. negative density estimates when using higher-order kernels) and computational burden. It is, therefore, worth stressing that here combining probit transformation and local log-quadratic density estimation in the (S, T) -domain achieves that bias reduction with no real extra complications compared to other estimators and fixes the boundary bias in an automatic way. Furthermore, these improvements are visible even in moderately large sample size, as the simulation study in Section 5 will show. The most obvious and practically relevant effect of this bias reduction is that a larger bandwidth can be used without oversmoothing. This results in smoother estimates, visually more pleasant. This is clear in Figure 3.2 (right panel), where $\tilde{c}^{(\tau,2)}$ is shown for the same data set as previously. Again, the bandwidth matrix in (2.9) was chosen via the cross-validation method suggested in Section 4.

3.3 Improved probit-transformation kernel copula density estimators with k -NN bandwidth

Theorems 3.2 and 3.3 reveal that the combination of probit transformation and local likelihood methods mostly cures the boundary bias problems for kernel copula density estimation. However, the fact remains that the suggested estimators have a variance behaving like

$$\text{Var}(\tilde{c}^{(\tau,p)}(u, v)) = C_p \frac{c(u, v)}{4\pi nh^2\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} + o((nh^2)^{-1}), \quad (3.17)$$

where $C_1 = 1$ and $C_2 = 5/2$, as $n \rightarrow \infty$. Hence, $\mathbb{V}\text{ar}(\tilde{c}^{(\tau,p)}(u, v))$ tends to grow unboundedly when (u, v) approaches one of the boundaries. Note that this is also the case for other copula density estimators attempting to correct the boundary bias, see for instance (Blumentritt, 2011, Chapter 4) and Janssen et al (2013) who obtain similar unbounded boundary variance for the Beta kernel and the Bernstein estimators. Facing the same situation in the univariate case, Geenens (2014) suggested to use k -Nearest-Neighbor (k -NN) type bandwidth in the transformed domain. Although barely used for standard kernel density estimation, k -NN bandwidths appeared totally appropriate in Geenens (2014)'s framework and, indeed, managed to stabilise the variance towards the boundaries. This is also the case in this setting as can be understood heuristically as follows.

Again, assume that the smoothing matrix \mathbf{H}_{ST} in (2.9) is diagonal, but instead of taking $\mathbf{H}_{ST} = h^2\mathbf{I}$ for some fixed value h , take a local smoothing matrix defined as $\mathcal{H}_{ST}^{(k)}(s, t) = D_k^2(s, t)\mathbf{I}$, where $D_k(s, t)$ is the Euclidean distance between (s, t) and the k th closest observation out of the sample (2.5) in \mathbb{R}^2 . Now it is k , or equivalently $\alpha = k/n$, that will play the role of the smoothing parameter in lieu of h . If \mathbf{K} had a compact support, α would be the proportion of observations actively entering the estimation of f_{ST} at any (s, t) – the interpretation roughly holds for the Gaussian kernel as well. Of course, $D_k(s, t)$ depends on the sample and is a random quantity. Along the same lines as in Mack and Rosenblatt (1979), one can show that $\mathbb{E}(1/D_k(s, t)) \simeq \frac{\pi f_{ST}(s, t)}{\alpha}$ and, together with $\mathbb{V}\text{ar}(\tilde{f}_{ST}^{(p)}(s, t)|D_k(s, t)) \simeq C_p \frac{f_{ST}(s, t)}{4\pi n D_k(s, t)}$, that

$$\mathbb{V}\text{ar}(\tilde{f}_{ST}^{(p)}(s, t)) \simeq C_p \frac{f_{ST}^2(s, t)}{4n\alpha}.$$

Now, through (2.3), one directly gets, for all $(u, v) \in (0, 1)^2$,

$$\mathbb{V}\text{ar}(\tilde{c}^{(\tau,p)}(u, v)) \simeq C_p \frac{c^2(u, v)}{4n\alpha}.$$

Unlike (3.17), this is no more proportional to $1/\{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))\}$ which grows unboundedly towards boundaries. This results in estimates more stable and much smoother towards the borders of \mathcal{I} .

In fact, given that the (long) tails of $\tilde{f}_{ST}^{(p)}$ in the transformed domain becomes the (short) boundary regions of $\tilde{c}^{(\tau,p)}$ in \mathcal{I} through the compressing back-transformation ($u = \Phi(s), v = \Phi(t)$), $\tilde{f}_{ST}^{(p)}$ must have *very smooth* tails in \mathbb{R}^2 to produce suitably smooth boundary behaviour for $\tilde{c}^{(\tau,p)}$. This is exactly what is achieved by using a k -NN bandwidth in the (S, T) -domain: local likelihood density estimators using k -NN bandwidth are, indeed, known to produce smoother estimates in the tails than their fixed-bandwidth counterparts, avoiding the occurrence of ‘spurious bumps’. Hence the appropriateness of the method here.




Figure 3.3: True Gaussian copula density with $\rho = 0.3$ (left), its local log-linear (middle) and log-quadratic (right) improved probit-transformation kernel estimators from the same sample ($n = 500$) as in Figure 3.1. Both estimates use an unconstrained bandwidth matrix of type k -NN chosen by cross-validation in the (S, T) -domain, see Section 4.

This is illustrated in Figure 3.3. Again, the previous simulated data set was used to produce the two ‘improved’ probit-transformation kernel copula density estimates shown in the middle (local log-linear) and the right panel (local log-quadratic), but this time a k -NN-type unconstrained bandwidth matrix $\mathcal{H}_{ST}^{(k)}(s, t)$ was used (see Section 4 for details). Compared to Figure 3.2, the estimates are much smoother along the boundaries now. For instance, using a k -NN-bandwidth mostly corrects the kink previously observed in the $(1, 0)$ corner. It is particularly clear for $\tilde{c}^{(\tau, 2)}$. The value of α selected for the case $p = 1$ was 0.1871, that for the case $p = 2$ was 0.4976. Again, the bias order reduction implied by local log-quadratic modelling allows a larger smoothing parameter to be used. As a result, this estimate $\tilde{c}^{(\tau, 2)}$ (right panel) has a smooth and visually pleasant appearance, but without oversmoothing. In fact, it is barely distinguishable from the true copula density (left panel). It happens that the estimator $\tilde{c}^{(\tau, 2)}$, when used in conjunction with a k -NN-type bandwidth, is strikingly good at recovering the shape of the underlying copula density while maintaining a visually pleasant amount of smoothness, see also Section 6.

4 Bandwidth choice

The behaviour of kernel estimators is known to be crucially dependent on their smoothing parameter, whose choice in practice is unanimously recognised as a very difficult problem, especially in more than one dimension. Here an effective way for selecting a suitable bandwidth matrix

$$\mathbf{H}_{ST} = \begin{pmatrix} h_1^2 & h_{12} \\ h_{12} & h_2^2 \end{pmatrix}$$

in (2.9) is suggested. It can be understood that the diagonal elements h_1^2 and h_2^2 of \mathbf{H}_{ST} quantify the amount of smoothing applied in the directions of the main s - and t -axes, hence their values will drive the overall smoothness of the resulting estimate $\tilde{f}_{ST}^{(p)}$ and eventually that of $\tilde{c}^{(\tau,p)}$. On the other hand, h_{12} sets the direction along which that smoothing mostly takes place in \mathbb{R}^2 . For instance, if \mathbf{K} is the bivariate Gaussian kernel, the local weights around $(s, t) \in \mathbb{R}^2$ will be set by the elliptical contour lines of the $\mathcal{N}_2((s, t)^t, \mathbf{H}_{ST})$ -distribution. If f_{ST} is stretched along a particular direction of \mathbb{R}^2 , which will be the case if c itself is so on the unit square, it is greatly beneficial to the estimator that smoothing be applied in that direction (Duong and Hazelton, 2005), and so h_{12} should be selected accordingly. If this is not the case, in particular if S and T are uncorrelated, then h_{12} may be set to 0. This motivates to separate the problem of selecting h_1 and h_2 from that of selecting h_{12} . The idea developed here looks for achieving this, in a way close in spirit to pre-sphering the observations (Wand and Jones, 1995, Section 4.6).

Consider the principal components decomposition of the $(n \times 2)$ -‘data matrix’ $\Xi \doteq (\hat{S}_i, \hat{T}_i)_{i=1}^n$. By construction, the \hat{S}_i ’s and \hat{T}_i ’s are centred, hence \hat{Q}_i and \hat{R}_i , the score of the i th observation on the first and second principal components, are given by

$$\hat{Q}_i = W_{11}\hat{S}_i + W_{12}\hat{T}_i, \quad \hat{R}_i = W_{21}\hat{S}_i + W_{22}\hat{T}_i, \quad (4.1)$$

where $W_1 = (W_{11}, W_{12})^t$ and $W_2 = (W_{21}, W_{22})^t$ are the eigenvectors of $\Xi^T \Xi$. Given that the transformation

$$\begin{pmatrix} Q \\ R \end{pmatrix} = \begin{pmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{pmatrix} \begin{pmatrix} S \\ T \end{pmatrix} \doteq \mathbf{W} \begin{pmatrix} S \\ T \end{pmatrix} \quad (4.2)$$

is only a linear reparametrisation of \mathbb{R}^2 , an estimate of f_{ST} can be readily obtained from an estimate of the density of (Q, R) , say f_{QR} . In addition, it is well known that the samples $\{\hat{Q}_i\}$ and $\{\hat{R}_i\}$ are uncorrelated, hence estimating f_{QR} via any kernel method from the sample $\{(\hat{Q}_i, \hat{R}_i); i = 1, \dots, n\}$ can be based on a diagonal bandwidth matrix $\mathbf{H}_{QR} = \text{diag}(h_Q^2, h_R^2)$ with little side effect. An idea is then to select h_Q and h_R independently via univariate procedures. Denote $\tilde{f}_Q^{(p)}$ and $\tilde{f}_R^{(p)}$ ($p = 1, 2$), the local log-polynomial estimators for the density of Q and R , respectively, based on the samples $\{\hat{Q}_i\}$ and $\{\hat{R}_i\}$ (see equations (6) and (7) in Loader (1996)). Of course, $\tilde{f}_Q^{(p)}$ only depends on one bandwidth h_Q and $\tilde{f}_R^{(p)}$ only depends on another bandwidth h_R . Then, h_Q can be selected via cross-validation (Loader, 1999, Section 5.3.3) as

$$h_Q = \arg \min_{h>0} \left\{ \int_{-\infty}^{\infty} \left\{ \tilde{f}_Q^{(p)}(q) \right\}^2 dq - \frac{2}{n} \sum_{i=1}^n \tilde{f}_{Q(-i)}^{(p)}(\hat{Q}_i) \right\}, \quad (4.3)$$

where, as usual in cross-validation procedures, $\tilde{f}_{Q(-i)}^{(p)}$ is the ‘leave-one-out’ version of $\tilde{f}_Q^{(p)}$ computed on all the observations but \hat{Q}_i . The value of h_R can be found similarly, and h_Q and h_R can be plugged into \mathbf{H}_{QR} for proceeding to bivariate estimation. However, optimal bandwidths for univariate density estimation are usually smaller than those for bivariate density estimation of f_{QR} . For the case $p = 1$ (local log-linear estimator), the asymptotic optimal bandwidth order is $n^{-1/5}$ for univariate density estimation and $n^{-1/6}$ in two dimensions. For the case $p = 2$ (local log-quadratic estimator), the asymptotic optimal bandwidth order is $n^{-1/9}$ for univariate density estimation and $n^{-1/10}$ in two dimensions. Hence, a fair choice for the bandwidth matrix for estimating f_{QR} seems to be

$$\mathbf{H}_{QR} = K_n^{(p)} \begin{pmatrix} h_Q^2 & 0 \\ 0 & h_R^2 \end{pmatrix},$$

with h_Q and h_R the two bandwidths found above by (univariate) cross-validation, and $K_n^{(1)} = n^{1/15}$ in the local log-linear case and $K_n^{(2)} = n^{1/45}$ in the local log-quadratic case. The estimate of f_{ST} can finally be obtained by linear back-transformation of the estimate of f_{QR} from the (Q, R) -domain to the (S, T) domain. It must be noted, though, that, again due to (4.2), this exactly amounts to directly estimating f_{ST} from $\{(\hat{S}_i, \hat{T}_i); i = 1, \dots, n\}$ using the bandwidth matrix

$$\mathbf{H}_{ST} = \mathbf{W}^{-1} \begin{pmatrix} h_Q^2 & 0 \\ 0 & h_R^2 \end{pmatrix} \mathbf{W}^{-1}.$$

When the estimator is to be built on a k -NN-type bandwidth matrix, the procedure is very similar. The data are transformed into the sample $\{(\hat{Q}_i, \hat{R}_i); i = 1, \dots, n\}$ via (4.1). Then, a suitable value of α in the Q -direction is computed as

$$\alpha_Q = \arg \min_{\alpha \in (0,1)} \left\{ \int_{-\infty}^{\infty} \left\{ \tilde{f}_Q^{(p)}(q) \right\}^2 dq - \frac{2}{n} \sum_{i=1}^n \tilde{f}_{Q(-i)}^{(p)}(\hat{Q}_i) \right\}, \quad (4.4)$$

i.e. exactly as in (4.3) but this time $\tilde{f}_Q^{(p)}$ is the estimator based on the k -NN bandwidth α . The value α_R is computed in the same way. Denote $\kappa = \alpha_Q/\alpha_R$. Then, the squared norm of a vector in the (Q, R) -domain will be computed as

$$\|(q, r)\|^2 = q^2 + \kappa^2 r^2. \quad (4.5)$$

The factor κ^2 naturally adjusts, through the obtained close-to-optimal values of α_Q and α_R , for the potential discrepancy in local geometry in the q - and r -directions. The bivariate estimation of f_{QR} at any

$(q, r) \in \mathbb{R}^2$ is carried out using the $k = K_n^{(p)} \times \alpha_Q \times n$ nearest neighbours of (q, r) , these being determined by the above distance. Here, $K_n^{(1)} = n^{-2/15}$ and $K_n^{(2)} = n^{-4/45}$, again for accounting for the difference in optimal α -orders in one and two dimensions. Finally, the estimate of f_{ST} is obtained by inverse linear transformation or, as set out in the fixed-bandwidth case, directly from the sample $\{(\hat{S}_i, \hat{T}_i); i = 1, \dots, n\}$ using an appropriate Mahalanobis-like distance. The main difference is that here, the employed distance makes use, through κ in (4.5), of relevant information in terms of optimal smoothing, not only in terms of the covariance structure of $\{(\hat{S}_i, \hat{T}_i); i = 1, \dots, n\}$ like the usual Mahalanobis distance. In this setting, the ‘smoothing parameters’ vector is, therefore, (α_Q, κ) .

It is acknowledged that this procedure may lack of sufficient theoretical support. For instance, it is known that pre-sphering the observations in the process of selecting the bandwidth matrix is justified only if the underlying density f_{ST} is bivariate normal, that is, in this framework, if c is the Gaussian copula. Likewise, choosing h_Q and h_R (or α_Q and α_R) independently via univariate procedures would be suitable in theory only if Q and R were independent, not only uncorrelated. The need for a correcting factor $K_n^{(p)}$ in the above bandwidth expressions may also seem like nothing less than a heuristic, *ad-hoc* correction. Having said this, it was found that this way of doing gave very reliable results, as illustrated in Figures 3.2 and 3.3. This will be even more obvious in the simulation study and the real data analysis detailed in the next sections. In addition, the suggested procedure, based on a twofold univariate cross-validation optimisation problem, is more stable numerically than one based on optimising a full, bivariate cross-validation criterion. This technique seems, therefore, an acceptable choice for selecting the bandwidth matrix in practice.

5 Simulation study

Here, Monte Carlo simulations results are presented to compare the practical behaviour of the probit-transformation estimators with that of their main competitors. All computations have been carried out using the R software and its freely available packages. Specifically, 12 estimators were considered:

- the ‘mirror reflection’ estimator, denoted $\hat{c}^{(m)}$ below, as suggested in Gijbels and Mielniczuk (1990). It was the first attempt at nonparametric copula density estimation, and remains a common choice for (ostensibly) correcting boundary bias. It will, therefore, be taken as benchmark. A first bandwidth matrix was obtained from the ‘augmented’ data set (made up of $9n$ ‘observations’ spread over an area 9 times bigger than \mathcal{I}) via the Normal reference rule, then the final matrix was obtained by multiplying the former by $(1/9)^{2/3} \simeq 0.23$ for adjusting for the effective sample size and range;
- the ‘naive’ probit-transformation estimator $\hat{c}^{(\tau)}$ (2.7) and its amended version $\hat{c}^{(\tau_{am})}$, whose idea is

exposed in Section 3.1 (for a general, non-diagonal matrix \mathbf{H}_{ST} , the amendment takes a slightly more complicated form than (3.11)). The bandwidth matrix \mathbf{H}_{ST} was selected via a direct plug-in method (Duong and Hazelton, 2003) in the transformed domain (S, T) ;

- the improved probit-transformation estimators $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$, given by (2.10), based on a k -NN-type bandwidth matrix selected via cross-validation as described at the end of Section 4. As already observed in Geenens (2014) in the univariate case, when based on a fixed-bandwidth matrix these estimators performed a little less well, so the results are not shown here. The optimisation problems (2.9) (local log-polynomial estimation of f_{ST}) and (4.4) (k -NN bandwidth selection) were solved using the relevant functions of the R package `locfit`. A R package directly implementing these improved probit-transformation estimators is in preparation;
- the Beta kernel estimator (Charpentier et al, 2007), denoted $\hat{c}^{(\beta)}$, with Chen (1999)'s further bias correction. Two smoothing parameters were considered: $h = 0.02$ ($\hat{c}_1^{(\beta)}$) and $h = 0.05$ ($\hat{c}_2^{(\beta)}$);
- the Bernstein copula density estimator (Bouezmarni et al, 2010, 2013, Janssen et al, 2013), denoted $\hat{c}^{(B)}$. Two smoothing parameters were considered: $k = 15$ ($\hat{c}_1^{(B)}$) and $k = 30$ ($\hat{c}_2^{(B)}$);
- the penalised hierarchical B -splines estimator (Kauermann et al, 2013), denoted $\hat{c}^{(p)}$, computed by the function `pencopula` in the eponymous R package. The vector of penalty coefficients was set to $\lambda = (10, 10)$ ($\hat{c}_1^{(p)}$), $\lambda = (100, 100)$ ($\hat{c}_2^{(p)}$), and $\lambda = (1000, 1000)$ ($\hat{c}_3^{(p)}$). The parameters d and D were set to 4 and 8, according to Kauermann et al (2013)'s simulations study.

$M = 1,000$ independent random samples $\{(U_i, V_i); i = 1, \dots, n\}$ of sizes $n = 200$, $n = 500$ and $n = 1000$ were generated from each of the following copulas:

- the independence copula (i.e., U_i 's and V_i 's drawn independently);
- the Gaussian copula, with parameters $\rho = 0.31$, $\rho = 0.59$ and $\rho = 0.81$;
- the Student t -copula with 10 degrees of freedom, with parameters $\rho = 0.31$, $\rho = 0.59$ and $\rho = 0.81$;
- the Student t -copula with 4 degrees of freedom, with parameters $\rho = 0.31$, $\rho = 0.59$ and $\rho = 0.81$;
- the Frank copula, with parameter $\theta = 1.86$, $\theta = 4.16$ and $\theta = 7.93$;
- the Gumbel copula, with parameter $\theta = 1.25$, $\theta = 1.67$ and $\theta = 2.5$;
- the Clayton copula, with parameter $\theta = 0.5$, $\theta = 1.67$ and $\theta = 2.5$.

For each family of copulas, the considered three values of the parameter roughly correspond to Kendall's τ 's equal to 0.2, 0.4 and 0.6, respectively. Of course, all the estimations only made use of the pseudo-

observations, i.e. the normalised ranks of the observations in the initially generated samples $\{U_i; i = 1, \dots, n\}$ and $\{V_i; i = 1, \dots, n\}$.

In order to assess the quality of the fit of an estimator \hat{c} for a given copula density c , the Mean Integrated L^2 -Error $\mathbb{E} \left(\iint_{\mathcal{I}} (\hat{c}(u, v) - c(u, v))^2 dudv \right)$ was estimated by the average over the $M = 1,000$ Monte Carlo replications of

$$ISE(\hat{c}) \simeq \frac{1}{(N+1)^2} \sum_{k_1=1}^N \sum_{k_2=1}^N \left\{ \hat{c} \left(\frac{k_1}{N+1}, \frac{k_2}{N+1} \right) - c \left(\frac{k_1}{N+1}, \frac{k_2}{N+1} \right) \right\}^2$$

with $N = 64$. The approximated MISE can be found in Tables 5.1, 5.2 and 5.3 for the three considered sample sizes. Note that, for ease of reading and interpretation, all the values are relative to the (approximated) MISE of the benchmark mirror estimator $\hat{c}^{(m)}$. For reference, the effective MISE of $\hat{c}^{(m)}$ is reported in italics in the last column of the table (which is, therefore, not on the same scale as the other values).

$n = 200$	$\hat{c}^{(\tau)}$	$\hat{c}^{(\tau_{am})}$	$\tilde{c}^{(\tau,1)}$	$\tilde{c}^{(\tau,2)}$	$\hat{c}_1^{(\beta)}$	$\hat{c}_2^{(\beta)}$	$\hat{c}_1^{(B)}$	$\hat{c}_2^{(B)}$	$\hat{c}_p^{(1)}$	$\hat{c}_p^{(2)}$	$\hat{c}_p^{(3)}$	$\hat{c}^{(m)}$
Indep	3.63	2.10	2.41	1.39	10.19	20.94	3.24	6.39	1.53	0.50	0.32	<i>0.02</i>
Gauss2	2.52	1.46	1.55	0.92	6.81	13.55	2.18	4.14	1.01	0.53	0.49	<i>0.03</i>
Gauss4	1.17	0.67	0.56	0.31	2.57	4.64	0.99	1.44	0.64	0.94	1.92	<i>0.08</i>
Gauss6	0.50	0.30	0.16	0.08	0.88	1.15	0.69	0.53	0.76	1.25	2.16	<i>0.37</i>
Std(10)2	2.05	1.15	1.28	0.72	5.30	10.95	1.71	3.18	0.96	0.77	0.90	<i>0.03</i>
Std(10)4	0.76	0.53	0.42	0.24	1.90	3.22	0.87	1.08	0.71	1.04	1.77	<i>0.12</i>
Std(10)6	0.33	0.28	0.13	0.10	0.78	0.85	0.68	0.48	0.82	1.22	1.89	<i>0.51</i>
Std(4)2	1.12	0.76	0.81	0.57	2.83	5.74	1.12	1.74	0.87	1.00	1.27	<i>0.07</i>
Std(4)4	0.44	0.41	0.32	0.25	1.23	1.79	0.76	0.71	0.79	1.12	1.59	<i>0.22</i>
Std(4)6	0.19	0.28	0.15	0.16	0.74	0.60	0.73	0.51	0.88	1.17	1.60	<i>0.82</i>
Frank2	3.54	2.00	2.17	1.28	9.06	18.22	2.81	5.53	1.20	0.37	0.30	<i>0.02</i>
Frank4	2.74	1.41	1.28	0.88	5.40	10.41	1.83	3.14	0.55	0.85	2.99	<i>0.03</i>
Frank6	1.31	0.62	0.50	0.51	1.73	2.92	1.05	1.09	0.58	1.66	4.03	<i>0.13</i>
Gumbel2	1.14	0.79	0.82	0.57	3.22	6.16	1.24	1.92	0.91	0.92	1.06	<i>0.06</i>
Gumbel4	0.36	0.42	0.29	0.30	1.18	1.52	0.77	0.71	0.84	1.08	1.46	<i>0.26</i>
Gumbel6	0.18	0.34	0.18	0.26	0.79	0.56	0.78	0.59	0.91	1.12	1.45	<i>1.08</i>
Clayton2	1.13	0.76	0.72	0.50	3.16	6.54	1.16	1.87	0.89	0.96	1.25	<i>0.06</i>
Clayton4	0.22	0.42	0.23	0.34	0.86	0.67	0.79	0.62	0.92	1.09	1.33	<i>0.75</i>
Clayton6	0.19	0.43	0.22	0.32	0.82	0.51	0.82	0.66	0.95	1.08	1.26	<i>1.77</i>

Table 5.1: (approximated) MISE relative to the MISE of the mirror-reflection estimator (last column), $n = 200$. Bold values show the minimum MISE for the corresponding copula (non-significantly different values are highlighted as well).

It turns out that the estimators $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$ are clearly the best, overall, on this L_2 -error criterion, and this for all sample sizes. They always dramatically improve on the Beta and Bernstein estimators, and they also do much better than the mirror reflection and the penalised B-splines estimators when the dependence is not close to null. When the dependence is very low, $\hat{c}^{(m)}$ and $\hat{c}^{(p)}$ do better, which can be easily

$n = 500$	$\hat{c}^{(\tau)}$	$\hat{c}^{(\text{ram})}$	$\tilde{c}^{(\tau,1)}$	$\tilde{c}^{(\tau,2)}$	$\hat{c}_1^{(\beta)}$	$\hat{c}_2^{(\beta)}$	$\hat{c}_1^{(B)}$	$\hat{c}_2^{(B)}$	$\hat{c}_p^{(1)}$	$\hat{c}_p^{(2)}$	$\hat{c}_p^{(3)}$	$\hat{c}^{(m)}$
Indep	3.37	2.31	2.54	1.27	7.90	13.99	2.06	4.21	1.53	0.51	0.23	0.01
Gauss2	2.22	1.47	1.63	0.78	5.24	8.82	1.41	2.59	0.98	0.63	0.68	0.02
Gauss4	0.79	0.55	0.48	0.23	1.88	2.55	0.80	0.84	0.63	0.97	2.46	0.06
Gauss6	0.31	0.24	0.13	0.06	0.74	0.60	0.70	0.39	0.73	1.23	2.54	0.30
Std(10)2	1.65	1.11	1.17	0.62	3.70	6.21	1.19	1.90	0.93	0.85	1.19	0.02
Std(10)4	0.52	0.42	0.34	0.17	1.35	1.70	0.74	0.63	0.69	1.08	2.19	0.10
Std(10)6	0.21	0.21	0.10	0.06	0.69	0.41	0.73	0.41	0.80	1.21	2.15	0.44
Std(4)2	0.78	0.64	0.60	0.46	1.88	2.86	0.86	0.96	0.81	1.01	1.58	0.05
Std(4)4	0.28	0.33	0.21	0.18	0.96	0.88	0.72	0.50	0.77	1.12	1.87	0.18
Std(4)6	0.12	0.22	0.10	0.11	0.70	0.31	0.77	0.48	0.87	1.17	1.77	0.74
Frank2	3.29	2.17	2.32	1.20	7.49	12.62	2.05	3.82	1.24	0.41	0.38	0.01
Frank4	2.49	1.41	1.40	0.91	4.39	6.55	1.53	2.11	0.58	0.80	4.50	0.02
Frank6	1.02	0.54	0.43	0.43	1.44	1.71	1.13	0.81	0.49	1.62	5.67	0.09
Gumbel2	0.83	0.71	0.65	0.47	2.16	3.20	0.90	1.09	0.87	0.98	1.30	0.05
Gumbel4	0.25	0.35	0.21	0.23	0.94	0.72	0.76	0.53	0.82	1.09	1.64	0.23
Gumbel6	0.11	0.26	0.12	0.18	0.77	0.36	0.82	0.57	0.92	1.12	1.56	0.99
Clayton2	0.85	0.67	0.61	0.40	2.20	3.34	0.88	1.06	0.84	1.02	1.57	0.04
Clayton4	0.15	0.32	0.14	0.21	0.79	0.37	0.79	0.56	0.91	1.09	1.43	0.69
Clayton6	0.15	0.35	0.13	0.19	0.81	0.40	0.85	0.65	0.95	1.08	1.32	1.67

Table 5.2: (approximated) MISE relative to the MISE of the mirror-reflection estimator (last column), $n = 500$. Bold values show the minimum MISE for the corresponding copula (non-significantly different values are highlighted as well).

understood. It is well known that the mirror reflection estimator efficiently deals with boundary effects only when the partial derivatives of c are 0 there (‘shoulder’). It is, therefore, particularly appropriate for the independence copula ($c \equiv 1$) and other very flat copula densities such as Gaussian or Frank with low dependence. The penalised B-splines estimator does even better when using a huge penalty for roughness, for obvious reasons. In all other cases, and particularly when the copula density becomes unbounded in some corners (but not only), $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$ dramatically outperform their competitors. In fact, mirror reflection and splines are not appropriate methods for estimating unbounded copula densities, and this is a real problem given that those are the most interesting cases in practice. By construction, the Beta kernel estimator always tends to be zero along boundaries (see for instance Figure 6.2 below), hence its even worse performance in this framework. The Bernstein estimator does better than $\hat{c}^{(\beta)}$, but cannot really compete with $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$. Of course, one can argue that the smoothing parameters used for $\hat{c}^{(\beta)}$, $\hat{c}^{(B)}$ and $\hat{c}^{(p)}$ have been selected mostly arbitrarily and are not adequate. This may be true, however, there is no simple, data-driven way of selecting those parameters, hence the choice was made subjectively exactly as a practitioner should have resolved to act. In addition, the above observations support that bad smoothing parameter choice is not the only reason for the poor performance of some of those estimators. Other evidence of that will be given in the next section on a real data set.

$n = 1000$	$\hat{c}^{(\tau)}$	$\hat{c}^{(\tau\text{am})}$	$\tilde{c}^{(\tau,1)}$	$\tilde{c}^{(\tau,2)}$	$\hat{c}_1^{(\beta)}$	$\hat{c}_2^{(\beta)}$	$\hat{c}_1^{(B)}$	$\hat{c}_2^{(B)}$	$\hat{c}_p^{(1)}$	$\hat{c}_p^{(2)}$	$\hat{c}_p^{(3)}$	$\hat{c}^{(m)}$
Indep	3.57	2.80	2.89	1.40	7.96	11.65	1.69	3.43	1.62	0.50	0.14	0.01
Gauss2	2.03	1.52	1.60	0.76	4.63	6.06	1.10	1.82	0.98	0.66	0.89	0.01
Gauss4	0.63	0.49	0.44	0.21	1.72	1.60	0.75	0.58	0.62	0.99	2.93	0.05
Gauss6	0.21	0.20	0.11	0.05	0.74	0.33	0.77	0.37	0.72	1.21	2.83	0.26
Std(10)2	1.36	1.06	1.04	0.55	3.07	3.98	0.96	1.24	0.86	0.87	1.48	0.02
Std(10)4	0.41	0.37	0.28	0.15	1.22	1.00	0.74	0.46	0.68	1.08	2.51	0.08
Std(10)6	0.15	0.18	0.08	0.05	0.71	0.24	0.79	0.41	0.84	1.21	2.36	0.39
Std(4)2	0.61	0.56	0.50	0.40	1.57	1.80	0.78	0.67	0.75	1.01	1.88	0.04
Std(4)4	0.21	0.27	0.17	0.15	0.88	0.51	0.75	0.42	0.75	1.12	2.07	0.16
Std(4)6	0.09	0.17	0.08	0.09	0.70	0.19	0.82	0.47	0.90	1.17	1.90	0.67
Frank2	3.31	2.42	2.57	1.35	7.16	9.63	1.70	2.95	1.31	0.45	0.49	0.01
Frank4	2.35	1.45	1.51	0.99	4.42	4.89	1.49	1.65	0.60	0.72	6.14	0.01
Frank6	0.96	0.52	0.45	0.44	1.51	1.19	1.35	0.76	0.65	1.58	7.25	0.07
Gumbel2	0.65	0.62	0.56	0.43	1.77	1.97	0.82	0.75	0.83	1.03	1.52	0.04
Gumbel4	0.18	0.28	0.16	0.19	0.89	0.41	0.78	0.47	0.81	1.10	1.78	0.21
Gumbel6	0.09	0.21	0.10	0.15	0.78	0.29	0.85	0.58	0.94	1.12	1.63	0.93
Clayton2	0.63	0.60	0.51	0.34	1.78	1.99	0.78	0.70	0.79	1.04	1.79	0.04
Clayton4	0.11	0.26	0.10	0.15	0.79	0.27	0.83	0.56	0.90	1.10	1.50	0.65
Clayton6	0.11	0.28	0.08	0.15	0.82	0.35	0.88	0.67	0.96	1.09	1.36	1.61

Table 5.3: (approximated) MISE relative to the MISE of the mirror-reflection estimator (last column), $n = 1000$. Bold values show the minimum MISE for the corresponding copula (non-significantly different values are highlighted as well).

In general, the local log-quadratic estimator $\tilde{c}^{(\tau,2)}$ is doing better than the local log-linear $\tilde{c}^{(\tau,1)}$, which was expected from the theoretical results. A notable exception, though, is in presence of high tail dependence, i.e. when the copula density tends very quickly to ∞ at one of the corners of \mathcal{I} , such as for Clayton and Gumbel copulas with high Kendall's τ . In fact, the extra smoothness guaranteed by local log-quadratic estimation tends to prevent the estimator from growing too quickly in the corners, and this is thus slightly detrimental in those cases. The same comment holds true when comparing the naive estimator $\hat{c}^{(\tau)}$ to its amended version $\hat{c}^{(\tau\text{am})}$. Generally, $\hat{c}^{(\tau\text{am})}$ has lower MISE than $\hat{c}^{(\tau)}$, except in the above-mentioned cases of high tail dependence. Indeed, the amendment prevents the estimate from exploding, even when the true density does. In any case, these 'naive' versions cannot really match the performance of the 'improved' versions $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$ on MISE, not to mention that their visual appearance is by far less pleasant.

Finally, other criteria were considered for comparing the different estimators, such as L_1 -error on the square $\iint_{\mathcal{I}} |\hat{c}(u,v) - c(u,v)| dudv$, L_1 - and L_2 -error on the first diagonal ($u = v$) and on a side ($u = 0.01$) of \mathcal{I} , or L_1 - and L_2 -error at a given point in one of the corners of \mathcal{I} ($(u,v) = (0.01,0.01)$). These results are available on request from the authors. All show, to the same extent as above, the superiority of the improved probit-transformation estimators over their competitors.

6 Real data analysis

In this section the well-known ‘Loss-ALAE’ dataset, reporting the indemnity payment (X_i 's) and allocated loss adjustment expense (Y_i 's) associated to 1,500 losses from an insurance company, is considered. Analysed in Frees and Valdez (1998), Klugman and Parsa (1999) and Denuit et al (2006), this dataset has since then become a classic in the copula literature. In particular, Frees and Valdez (1998) mentioned that the Gumbel copula with $\hat{\theta} = 1.453$ provides an excellent fit. The data set initially contains 34 censored observations, that were excluded here as the suggested estimators were not designed to take censorship into account. Using more advanced model selection ideas, Chen et al (2010) also found that the Gumbel copula (with the same parameter $\hat{\theta}$) fits the dataset (restricted to its complete cases) the best out of most of the usual parametric copula models. The aim here is to test the probit-transformation estimators $\hat{c}^{(\tau,p)}$ ($p = 1, 2$) (and their competitors) against that parametric ‘gold standard’, shown in Figure 6.1 (up-left).

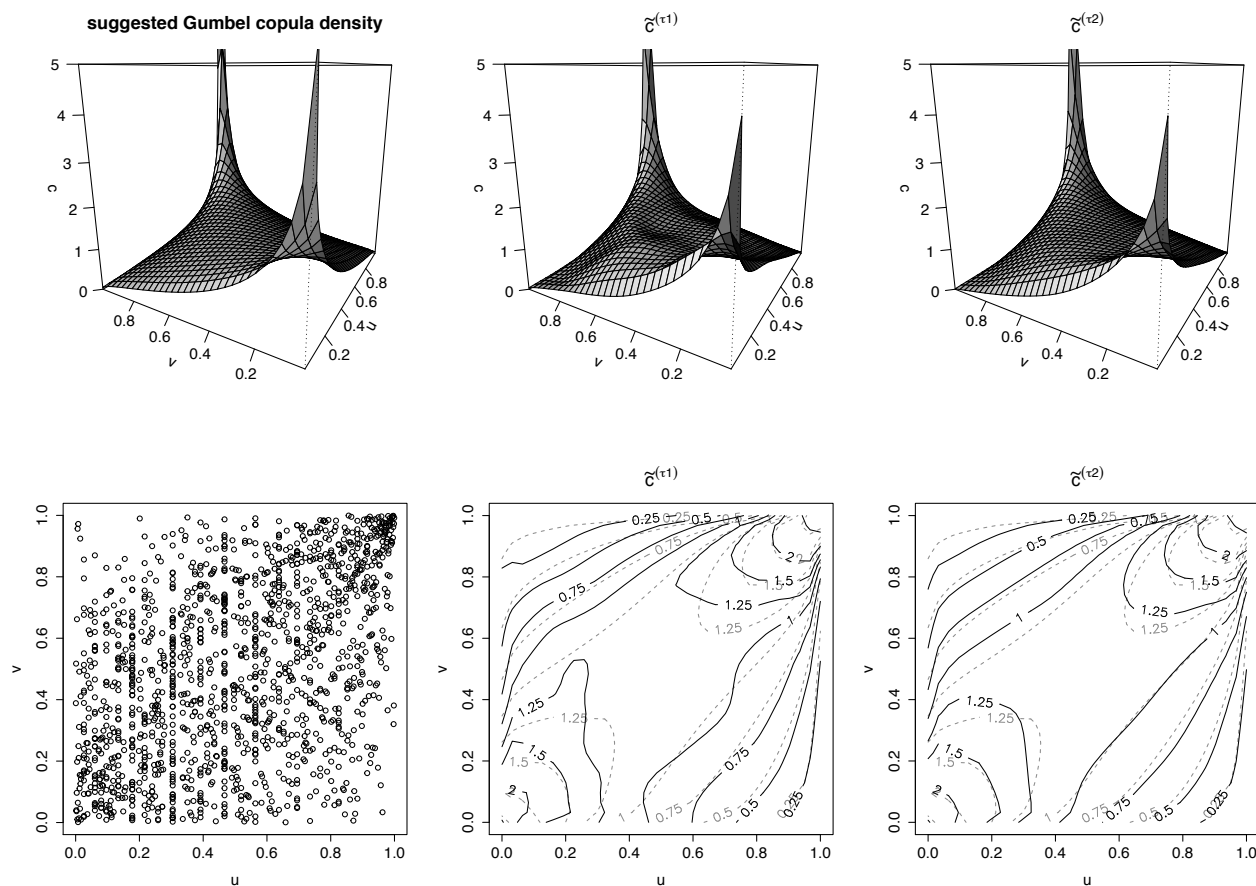


Figure 6.1: Loss-ALAE dataset: suggested parametric copula density (Gumbel with parameter $\hat{\theta} = 1.453$; upper-left panel) and probit-transformation estimates with $p = 1$ (middle column) and $p = 2$ (right column). The upper line shows 3-d views and the bottom line shows contour lines, superimposed on the Gumbel copula density contour lines. Pseudo-observations are shown in the bottom-left panel.

The two probit-transformation estimators (local log-linear and local log-quadratic) were first fit to the

data set. In both cases, a k -NN bandwidth matrix was used. Using the selection rule prescribed in Section 4, the parameters $(\alpha, \kappa) = (0.24, 1.28)$ for $p = 1$ and $(\alpha, \kappa) = (0.51, 1.01)$ for $p = 2$ were obtained in an automatic manner. The estimator $\tilde{c}^{(\tau,2)}$ is, again, very similar to the parametric fit. In particular, it has that very smooth and pleasant appearance of parametric estimates, while being based on a fully nonparametric procedure. Reproducing ‘parametric smoothness’ without sacrificing any flexibility is, of course, a huge achievement for $\tilde{c}^{(\tau,2)}$. Naturally, $\tilde{c}^{(\tau,1)}$ is less smooth (smaller value of α than for $\tilde{c}^{(\tau,2)}$, for the reasons explained at the end of Section 3.2), but is still totally acceptable. Both nonparametric estimates suggest that the true underlying copula density decays towards the $(0, 1)$ -corner quicker than what the Gumbel model shows (this is particularly clear from the contour lines). Admittedly, there is no way of knowing what is the truth here. However, $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$ are based only on the data (it is visually obvious that the upper-left corner of \mathcal{I} is much less endowed in data than the bottom-right corner), and not on any prior assumption. On the contrary, the Gumbel copula density is inherently symmetric in u and v . The peak in the density at $(0, 0)$ also appears less high on the nonparametric estimates than on the Gumbel copula density.

Figure 6.2 shows the competitors on the same data set: the mirror reflection estimator, two Beta kernel estimators, two Bernstein estimators and two penalised B -splines estimators. Of course, $\hat{c}^{(m)}$ cannot cope with this unbounded copula density. For the other three methods, producing an estimate reasonably smooth required a value of the smoothing parameter (h for Beta kernel estimators, k for Bernstein estimators and λ for penalised B -splines) preventing correct estimation of the peaks at $(0, 0)$ and $(1, 1)$. To get estimates showing a peak at $(1, 1)$ of roughly the right magnitude, one needed to use smoothing parameters producing unacceptably undersmoothed estimation elsewhere on \mathcal{I} , yet not even able to properly catch the peak at $(0, 0)$. If the Gumbel copula density is assumed to be close to the truth for this data set, then there is no question that $\tilde{c}^{(\tau,1)}$ and $\tilde{c}^{(\tau,2)}$ are, by far, the best. This, also, illustrates that the results obtained in the simulations are not only due to bad smoothing parameter choices.

7 Concluding remarks

Development of efficient kernel-type methods for nonparametric copula modelling have been delayed owing mainly to the bounded support of copulas, namely the unit square \mathcal{I} . It is, indeed, well known that kernel estimators heavily suffer from boundary bias issues, which are not trivial to fix. In this paper, a new kernel-type estimator for the copula density has been proposed. It is based on the probit-transformation idea suggested in Charpentier et al (2007) and studied in full in the univariate case in Geenens (2014). This ‘improved probit-transformation estimator’ deals with boundary bias in a very natural way. In addition,

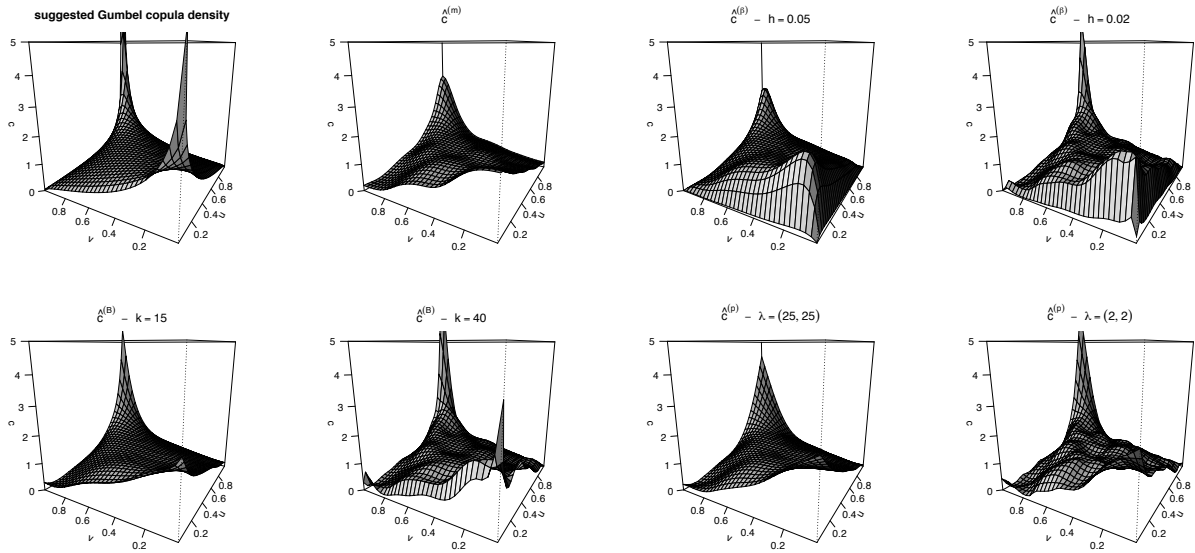


Figure 6.2: Loss-ALAE dataset: suggested Gumbel copula density ($\hat{\theta} = 1.456$), mirror reflection estimator, Beta kernel estimators with $h = 0.05$ and $h = 0.02$, Bernstein estimators with $k = 15$ and $k = 40$ and penalized B -splines estimators with $d = 4, D = 8$ and with $\lambda = (25, 25)$ and $\lambda = (2, 2)$

it has been seen to easily cope with potentially unbounded copula densities, which are the common and interesting cases in copula modelling. An easy-to-implement selection rule for the necessary smoothing parameters has also been proposed. This procedure has been seen to be very stable and to give very good results in practice. In particular, a version of the estimator ($\hat{c}^{(\tau, 2)}$ with k -NN-type bandwidth matrix) is able to reproduce the smooth and pleasant appearance of parametric models, while keeping the flexibility of fully nonparametric estimation procedures. A comprehensive simulation study has emphasised the very good practical performance of that estimator compared to its main competitors.

Several important points remain to be studied, though. First, as of now, the theoretical properties of the estimator have been derived under the assumption of i.i.d. sampling, making use of the strong approximation for the empirical copula process provided by Proposition 4.2 of Segers (2012). To the best of these authors' knowledge, this result has not been proved in the case of weakly dependent observations. It would be particularly significant to investigate this in a near future, given the predominant place recently found by copula modelling in the setting of time series, notably in finance. Other directions for future research would look for using the new copula density estimator in a variety of problems, for instance copula goodness-of-fit tests (Fermanian, 2005, Scaillet, 2007) or nonparametric conditional density estimation (Faugeras, 2009). Finally, it must be said that, in theory, the idea presented in this paper is not bound to the bivariate case but extends in a straightforward way to higher dimensional copulas as well. Of course, in practice, this is wise only within the limits allowed by the *curse of dimensionality*.

Acknowledgments

The first author was supported by a Faculty Research Grant from the Faculty of Science, University of New South Wales (Australia). The second author acknowledges additional funding provided by the Natural Sciences and Engineering Research Council of Canada. The third author was supported by an A.R.C. contract from the *Communauté Française de Belgique* and by the IAP research network grant nr. P7/06 of the Belgian government (Belgian Science Policy).

A Appendix

First a technical lemma, which may be of interest of its own, is stated.

Lemma A.1. *Under Assumptions 3.1-3.3, the density f_{ST} of the vector $(S, T) = (\Phi^{-1}(U), \Phi^{-1}(V))$ is uniformly bounded on \mathbb{R}^2 , and so are its partial derivatives up to the second order.*

Proof. From Assumption 3.3, one easily obtains that, for all $(u, v) \in (0, 1)^2$,

$$\begin{aligned}
c(u, v) &\leq K_{00} \min\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right) \\
&= K_{00} \min\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right)^\alpha \min\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right)^{1-\alpha} \quad \forall \alpha \in (0, 1) \\
&\leq K_{00} \min\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right)^\alpha \max\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right)^{1-\alpha} \\
&\leq K_{00} \min\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right)^\beta \max\left(\frac{1}{u(1-u)}, \frac{1}{v(1-v)}\right)^\beta \quad \text{with } \beta = \max(\alpha, 1-\alpha) \\
&= K_{00} \left(\frac{1}{uv(1-u)(1-v)}\right)^\beta, \quad \text{for some } \beta \in (1/2, 1). \tag{A.1}
\end{aligned}$$

In particular, when approaching the $(0, 0)$ -corner, the above implies that

$$\lim_{(u,v) \rightarrow (0,0)} u^\beta v^\beta c(u, v) = k_{00},$$

for some constant $0 \leq k_{00} < \infty$ (and similar towards the boundaries and the other corners of \mathcal{I}). Now, applying Theorem 1 of Lawlor (2012), one can show that this implies that, for $i, j = 0, 1, 2$ s.t. $i + j \leq 2$, there exist constants $k_{ij} < \infty$ such that

$$\lim_{(u,v) \rightarrow (0,0)} u^{\beta+i} v^{\beta+j} \left| \frac{\partial^{i+j} c(u, v)}{\partial u^i \partial v^j} \right| = k_{ij}$$

(and again, similar results hold at boundaries and at the other corners of \mathcal{I}). Given that $c(u, v)$ is assumed

to be twice continuously differentiable everywhere on the interior of \mathcal{I} (i.e., $c(u, v)$ and all its mixed partial derivatives of the first two orders can only go unbounded towards the boundaries of \mathcal{I}), this allows one to write that, there exist $\beta \in (1/2, 1)$ and bounded constants K_{ij} such that

$$\left| \frac{\partial^{i+j} c(u, v)}{\partial u^i \partial v^j} \right| \leq \frac{K_{ij}}{u^{\beta+i} (1-u)^{\beta+i} v^{\beta+j} (1-v)^{\beta+j}}, \quad (\text{A.2})$$

for all $(u, v) \in (0, 1)^2$. Then, (A.1) in (2.1) yields

$$f_{ST}(s, t) \leq \frac{K_{00} \phi(s) \phi(t)}{\Phi(s)^\beta (1-\Phi(s))^\beta \Phi(t)^\beta (1-\Phi(t))^\beta}$$

for all $(s, t) \in \mathbb{R}^2$. Given that, from the known properties of the normal distribution, $\sup_{s \in \mathbb{R}} \frac{\phi(s)}{\Phi(s)^\beta (1-\Phi(s))^\beta}$ is bounded for any constant $\beta < 1$, there exists a constant M such that $\sup_{(s, t) \in \mathbb{R}^2} f_{ST}(s, t) \leq M < \infty$.

Now from (3.8) one can write

$$\left| \frac{\partial f_{ST}}{\partial s}(s, t) \right| \leq \left| \frac{\partial c}{\partial u}(\Phi(s), \Phi(t)) \right| \phi^2(s) \phi(t) + |s| c(\Phi(s), \Phi(t)) \phi(s) \phi(t).$$

From (A.1) and (A.2) with $i = 1, j = 0$, one gets

$$\left| \frac{\partial f_{ST}}{\partial s}(s, t) \right| \leq \frac{K_{10} \phi^2(s) \phi(t)}{\Phi(s)^{\beta+1} (1-\Phi(s))^{\beta+1} \Phi(t)^\beta (1-\Phi(t))^\beta} + |s| \frac{K_{00} \phi(s) \phi(t)}{\Phi(s)^\beta (1-\Phi(s))^\beta \Phi(t)^\beta (1-\Phi(t))^\beta}.$$

Take $\gamma = (\beta + 1)/2 < 1$ and see that, as above, $\sup_{s \in \mathbb{R}} \frac{\phi(s)}{\Phi(s)^\gamma (1-\Phi(s))^\gamma}$ is bounded, so that the first term is uniformly bounded on \mathbb{R}^2 . The second term is uniformly bounded as well, as $\sup_{s \in \mathbb{R}} \frac{s \phi(s)}{\Phi(s)^\beta (1-\Phi(s))^\beta}$ is also bounded for any $\beta < 1$. All second-order partial derivatives of f_{ST} can be uniformly bounded in the exact same way using (A.2) in (3.9) and similar. \square

Proof of Proposition 3.1

Denote

$$C_n(u, v) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{U_i \leq u, V_i \leq v\}},$$

i.e. the ‘ideal’ version of the empirical copula (3.4) using genuine observations (U_i, V_i) ’s, and the corresponding empirical process $\{\mathbb{B}_n(u, v) : (u, v) \in \mathcal{I}\}$, with

$$\mathbb{B}_n(u, v) = \sqrt{n}(C_n(u, v) - C(u, v)).$$

Also, define the process $\{\mathbb{G}_n(u, v) : (u, v) \in \mathcal{I}\}$ with

$$\mathbb{G}_n(u, v) = \mathbb{B}_n(u, v) - \frac{\partial \mathcal{C}}{\partial u}(u, v)\mathbb{B}_n(u, 1) - \frac{\partial \mathcal{C}}{\partial v}(u, v)\mathbb{B}_n(1, v). \quad (\text{A.3})$$

Segers (2012) shows that, under Assumptions 3.1-3.3, the empirical copula process $\mathbb{C}_n(u, v)$ (see (3.7)) and $\mathbb{G}_n(u, v)$ are such that

$$\sup_{(u,v) \in \mathcal{I}} |\mathbb{C}_n(u, v) - \mathbb{G}_n(u, v)| = O_{\text{a.s.}} \left(n^{-1/4} (\log n)^{1/2} (\log \log n)^{1/4} \right) \quad n \rightarrow \infty. \quad (\text{A.4})$$

Now, see that

$$\sqrt{nh^2} \left(\hat{f}_{ST}(s, t) - \mathbb{E} \left(\hat{f}_{ST}^*(s, t) \right) \right) = \frac{1}{h} \iint_{\mathcal{I}} \phi \left(\frac{s - \Phi^{-1}(u)}{h} \right) \phi \left(\frac{t - \Phi^{-1}(v)}{h} \right) d\mathbb{C}_n(u, v).$$

Integrating by parts, as in the proof of Theorem 6 of Fermanian et al (2004), one gets

$$\begin{aligned} & \sqrt{nh^2} \left(\hat{f}_{ST}(s, t) - \mathbb{E} \left(\hat{f}_{ST}^*(s, t) \right) \right) \\ &= \frac{1}{h} \iint_{\mathcal{I}} \mathbb{C}_n(u, v) \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \phi' \left(\frac{t - \Phi^{-1}(v)}{h} \right) \frac{du}{h\phi(\Phi^{-1}(u))} \frac{dv}{h\phi(\Phi^{-1}(v))} \\ &= \frac{1}{h} \iint_{\mathcal{I}} \mathbb{G}_n(u, v) \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \phi' \left(\frac{t - \Phi^{-1}(v)}{h} \right) \frac{du}{h\phi(\Phi^{-1}(u))} \frac{dv}{h\phi(\Phi^{-1}(v))} \\ & \quad + R_n(s, t) \end{aligned} \quad (\text{A.5})$$

where

$$\begin{aligned} |R_n(s, t)| &\leq \frac{1}{h} \sup_{(u,v) \in \mathcal{I}} |\mathbb{C}_n(u, v) - \mathbb{G}_n(u, v)| \left\{ \int_0^1 \left| \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \right| \frac{du}{h\phi(\Phi^{-1}(u))} \right\}^2 \\ &= \frac{1}{h} \sup_{(u,v) \in \mathcal{I}} |\mathbb{C}_n(u, v) - \mathbb{G}_n(u, v)| \left\{ \int_{\mathbb{R}} |z| \phi(z) dz \right\}^2 \\ &= O_{\text{a.s.}} \left(n^{-1/4} h^{-1} (\log n)^{1/2} (\log \log n)^{1/4} \right) \\ &= o_{\text{a.s.}}(1), \end{aligned}$$

given that $\mathbb{E}(|\mathcal{N}(0, 1)|) < \infty$ and the conditions on the bandwidth h . Call

$$J_{st,h}(u, v) = \phi \left(\frac{s - \Phi^{-1}(u)}{h} \right) \phi \left(\frac{t - \Phi^{-1}(v)}{h} \right),$$

so that

$$dJ_{st,h}(u, v) = \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \phi' \left(\frac{t - \Phi^{-1}(v)}{h} \right) \frac{du}{h\phi(\Phi^{-1}(u))} \frac{dv}{h\phi(\Phi^{-1}(v))}.$$

Plugging (A.3) in (A.5) yields

$$\begin{aligned} \sqrt{nh^2} \left(\hat{f}_{ST}(s, t) - \mathbb{E} \left(\hat{f}_{ST}^*(s, t) \right) \right) &= \frac{1}{h} \iint_{\mathcal{I}} \mathbb{B}_n(u, v) dJ_{st,h}(u, v) \\ &\quad - \frac{1}{h} \iint_{\mathcal{I}} \frac{\partial C}{\partial u}(u, v) \mathbb{B}_n(u, 1) dJ_{st,h}(u, v) \\ &\quad - \frac{1}{h} \iint_{\mathcal{I}} \frac{\partial C}{\partial v}(u, v) \mathbb{B}_n(1, v) dJ_{st,h}(u, v) + R_n(s, t) \\ &\doteq A_n(s, t) + B_{n,1}(s, t) + B_{n,2}(s, t) + R_n(s, t). \end{aligned}$$

The process $\mathbb{B}_n(u, v)$ being the bivariate empirical process based on genuine observations, $A_n(s, t)$ is, in fact, $\sqrt{nh^2} \left(\hat{f}_{ST}^*(s, t) - \mathbb{E} \left(\hat{f}_{ST}^*(s, t) \right) \right)$, for which classical kernel smoothing theory states that

$$A_n(s, t) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \frac{f_{ST}(s, t)}{4\pi} \right).$$

The terms $B_{n,1}$ and $B_{n,2}$ can be worked out explicitly. This is done below for $B_{n,1}$ only ($B_{n,2}$ can be treated in the exact same way). Write

$$B_{n,1}(s, t) = -\frac{1}{h} \int_0^1 \mathbb{B}_n(u, 1) \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \left\{ \int_0^1 \frac{\partial C}{\partial u}(u, v) \phi' \left(\frac{t - \Phi^{-1}(v)}{h} \right) \frac{dv}{h\phi(\Phi^{-1}(v))} \right\} \frac{du}{h\phi(\Phi^{-1}(u))}$$

and proceed with

$$\begin{aligned} \psi(u) &\doteq \int_0^1 \frac{\partial C}{\partial u}(u, v) \phi' \left(\frac{t - \Phi^{-1}(v)}{h} \right) \frac{dv}{h\phi(\Phi^{-1}(v))} \\ &= \int_{\mathbb{R}} \frac{\partial C}{\partial u}(u, \Phi(t - hz)) \phi'(z) dz \end{aligned}$$

with the change of variable $z = h^{-1}(t - \Phi^{-1}(v))$. As $C(u, v) = F_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))$, this is also

$$\psi(u) = \frac{1}{\phi(\Phi^{-1}(u))} \int_{\mathbb{R}} \frac{\partial F_{ST}}{\partial s}(\Phi^{-1}(u), t - hz) \phi'(z) dz. \quad (\text{A.6})$$

Taylor-expanding, one gets

$$\begin{aligned} \frac{\partial F_{ST}}{\partial s}(\Phi^{-1}(u), t - hz) &= \frac{\partial F_{ST}}{\partial s}(\Phi^{-1}(u), t) - hz \frac{\partial^2 F_{ST}}{\partial s \partial t}(\Phi^{-1}(u), t) \\ &\quad + \frac{1}{2} h^2 z^2 \frac{\partial^3 F_{ST}}{\partial s \partial t^2}(\Phi^{-1}(u), t) - \frac{1}{6} h^3 z^3 \frac{\partial^4 F_{ST}}{\partial s \partial t^3}(\Phi^{-1}(u), t - h^* z), \end{aligned}$$

where $h^* \in (0, h)$, of which only odd powers of z will remain in (A.6) as ϕ' is an odd function. Hence,

$$\begin{aligned}\psi(u) &= -\frac{1}{\phi(\Phi^{-1}(u))} \left(hf_{ST}(\Phi^{-1}(u), t) \int_{\mathbb{R}} z\phi'(z) dz + \frac{1}{6}h^3 \int_{\mathbb{R}} z^3 \frac{\partial^2 f_{ST}}{\partial t^2}(\Phi^{-1}(u), t - h^*z)\phi'(z) dz \right) \\ &= \frac{1}{\phi(\Phi^{-1}(u))} \left(hf_{ST}(\Phi^{-1}(u), t) + \frac{1}{6}h^3\psi^*(u) \right),\end{aligned}$$

since $\int_{\mathbb{R}} z\phi'(z) dz = -1$, denoting $\psi^*(u) = -\int_{\mathbb{R}} z^3 \frac{\partial^2 f_{ST}}{\partial t^2}(\Phi^{-1}(u), t - h^*z)\phi'(z) dz$.

It follows that

$$\begin{aligned}B_{n,1}(s, t) &= -\frac{1}{h} \int_0^1 \mathbb{B}_n(u, 1) \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \psi(u) \frac{du}{h\phi(\Phi^{-1}(u))} \\ &= -\int_0^1 \mathbb{B}_n(u, 1) f_{ST}(\Phi^{-1}(u), t) \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \frac{du}{h\phi^2(\Phi^{-1}(u))} \\ &\quad - \frac{1}{6}h^2 \int_0^1 \mathbb{B}_n(u, 1) \phi' \left(\frac{s - \Phi^{-1}(u)}{h} \right) \psi^*(u) \frac{du}{h\phi^2(\Phi^{-1}(u))} \\ &\doteq B_{n,11}(s, t) + B_{n,12}(s, t).\end{aligned}$$

The change of variable $w = \frac{s - \Phi^{-1}(u)}{h}$ yields

$$\begin{aligned}B_{n,11}(s, t) &= \int_{\mathbb{R}} \mathbb{B}_n(\Phi(s - hw), 1) \phi'(w) \frac{f_{ST}(s - hw, t)}{\phi(s - hw)} dw \\ &= \mathbb{B}_n(\Phi(s), 1) \frac{f_{ST}(s, t)}{\phi(s)} \int_{\mathbb{R}} \phi'(w) dw \\ &\quad - \frac{f_{ST}(s, t)}{\phi(s)} \int_{\mathbb{R}} (\mathbb{B}_n(\Phi(s), 1) - \mathbb{B}_n(\Phi(s) - h\phi(s - h^*w), 1)) \phi'(w) dw \\ &\quad + h \int_{\mathbb{R}} \mathbb{B}_n(\Phi(s - hw), 1) \phi'(w) \frac{\partial}{\partial s} \left\{ \frac{f_{ST}(s - h^{**}w, t)}{\phi(s - h^{**}w)} \right\} dw,\end{aligned}$$

where h^*, h^{**} are both between 0 and h . The first term in this decomposition is 0, as ϕ' is an odd function.

The second term tends to 0 as $n \rightarrow \infty$. This is because $\mathbb{B}_n(\cdot, 1)$ is nothing else than a usual univariate uniform empirical process. Einmahl and Ruymgaart (1987) studied its modulus of continuity and their Theorem 3.1(b) allows one to write, under these assumptions,

$$\sup_{u, u' \in [0, 1], |u - u'| \leq h} |\mathbb{B}_n(u, 1) - \mathbb{B}_n(u', 1)| = \sqrt{2h \log h^{-1}} \quad \text{almost surely as } n \rightarrow \infty.$$

Hence the second term can readily be seen to be $O_{\text{a.s.}}(\sqrt{h \log h^{-1}})$ as $n \rightarrow \infty$, i.e. $o_{\text{a.s.}}(1)$ as $h \rightarrow 0$. The integral in the third term is $O_P(1)$. Given that $\frac{\partial}{\partial s} \left\{ \frac{f_{ST}(s - h^{**}w, t)}{\phi(s - h^{**}w)} \right\} = \frac{\partial f_{ST}/\partial s}{\phi}(s - h^{**}w, t) + (s - h^{**}w) \frac{f_{ST}}{\phi}(s -$

$h^{**}w, t)$, one can write

$$\begin{aligned} & \int_{\mathbb{R}} \phi'(w) \frac{\partial}{\partial s} \left\{ \frac{f_{ST}(s - h^{**}w, t)}{\phi(s - h^{**}w)} \right\} dw \\ &= - \int_{\mathbb{R}} w \frac{\phi(w)}{\phi(s - h^{**}w)} \frac{\partial f_{ST}}{\partial s}(s - h^{**}w, t) dw - \int_{\mathbb{R}} (s - h^{**}w) w \frac{\phi(w)}{\phi(s - h^{**}w)} f_{ST}(s - h^{**}w, t) dw. \end{aligned}$$

This is bounded for n large enough, because f_{ST} and $\frac{\partial f_{ST}}{\partial s}$ are uniformly bounded on \mathbb{R}^2 by Lemma A.1, and

$$\int_{\mathbb{R}} \left| w^j \frac{\phi(w)}{\phi(s - h^{**}w)} \right| dw < \infty, \quad j = 1, 2,$$

for $h^{**} < 1$, which eventually occurs as $n \rightarrow \infty$. As $\sup_w \mathbb{B}_n(\Phi(s - hw), 1) = O_P(1)$, one can see that

$$|B_{n,11}(s, t)| = O_P(h) + O_{\text{a.s.}}(\sqrt{h \log h^{-1}}) = o_P(1) \quad \text{as } n \rightarrow \infty.$$

The same can be written for $|B_{n,12}(s, t)|$, given that $\psi^*(u)$ is a uniformly bounded function of u on $[0, 1]$, from Lemma A.1. Hence,

$$\sqrt{nh^2} \left(\hat{f}_{ST}(s, t) - \mathbb{E} \left(\hat{f}_{ST}^*(s, t) \right) \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \frac{f_{ST}(s, t)}{4\pi} \right).$$

The final result follows by seeing that, using classical kernel density estimation results,

$$\mathbb{E} \left(\hat{f}_{ST}^*(s, t) \right) = f_{ST}(s, t) + \frac{1}{2} h^2 \left(\frac{\partial^2 f_{ST}}{\partial s^2}(s, t) + \frac{\partial^2 f_{ST}}{\partial t^2}(s, t) \right) + o(h^2) \quad \text{as } n \rightarrow \infty.$$

□

Proof of Theorem 3.1

Result (3.6) can be written at $(s, t) = (\Phi^{-1}(u), \Phi^{-1}(v))$:

$$\begin{aligned} & \sqrt{nh^2} \left(\hat{f}_{ST}(\Phi^{-1}(u), \Phi^{-1}(v)) - f_{ST}(\Phi^{-1}(u), \Phi^{-1}(v)) - h^2 b_{ST}(\Phi^{-1}(u), \Phi^{-1}(v)) \right) \\ & \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \sigma_{ST}^2(\Phi^{-1}(u), \Phi^{-1}(v)) \right), \end{aligned}$$

which implies

$$\sqrt{nh^2} \left(\frac{\hat{f}_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} - \frac{f_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} - h^2 \frac{b_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))} \right) \xrightarrow{\mathcal{L}} \mathcal{N} \left(0, \frac{\sigma_{ST}^2(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi^2(\Phi^{-1}(u))\phi^2(\Phi^{-1}(v))} \right). \quad (\text{A.7})$$

Now,

$$\frac{\sigma_{ST}^2(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi^2(\Phi^{-1}(u))\phi^2(\Phi^{-1}(v))} = \frac{f_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{4\pi\phi^2(\Phi^{-1}(u))\phi^2(\Phi^{-1}(v))} = \frac{c(u, v)}{4\pi\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))},$$

using (2.2). In addition, from (3.9), it is easily seen that $\frac{b_{ST}(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))}$ has the form (3.10). Using (2.2)/(2.3) in (A.7) then concludes the proof. \square

Proof of Theorem 3.2

From Proposition 3.1 one can conclude that

$$|\hat{f}_{ST}(s, t) - \hat{f}_{ST}^*(s, t)| = o_P((nh^2)^{-1/2}). \quad (\text{A.8})$$

Now, from (3.15), one has

$$|\tilde{f}_{ST}^{(1)}(s, t) - \tilde{f}_{ST}^{*(1)}(s, t)| = \left| \hat{f}_{ST}(s, t) \exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}(s, t) \right\} - \hat{f}_{ST}^*(s, t) \exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}^*(s, t) \right\} \right|,$$

where

$$\hat{\psi}_{ST}(s, t) = \left(\frac{\partial \hat{f}_{ST}(s, t) / \partial s}{\hat{f}_{ST}(s, t)} \right)^2 + \left(\frac{\partial \hat{f}_{ST}(s, t) / \partial t}{\hat{f}_{ST}(s, t)} \right)^2$$

and equivalently for the star version. Then,

$$\begin{aligned} |\tilde{f}_{ST}^{(1)}(s, t) - \tilde{f}_{ST}^{*(1)}(s, t)| &\leq |\hat{f}_{ST}(s, t) - \hat{f}_{ST}^*(s, t)| \exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}(s, t) \right\} \\ &\quad + \hat{f}_{ST}^*(s, t) \left| \exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}(s, t) \right\} - \exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}^*(s, t) \right\} \right|. \end{aligned}$$

The first term is $o_P((nh^2)^{-1/2})$, given (A.8) and $\exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}(s, t) \right\} \xrightarrow{P} 1$ as $n \rightarrow \infty$ (i.e. as $h^2 \rightarrow 0$).

The second term can be written

$$\frac{1}{2} h^2 \hat{f}_{ST}^*(s, t) \exp \left\{ -\frac{1}{2}h^2\hat{\psi}_{ST}^*(s, t) \right\} \left| \hat{\psi}_{ST}(s, t) - \hat{\psi}_{ST}^*(s, t) \right|, \quad (\text{A.9})$$

where $\check{\psi}_{ST}^*(s, t)$ is between $\hat{\psi}_{ST}(s, t)$ and $\hat{\psi}_{ST}^*(s, t)$. Now, it can be seen that

$$\frac{\left| \hat{\psi}_{ST}(s, t) - \hat{\psi}_{ST}^*(s, t) \right|}{\left| \frac{\partial \hat{f}_{ST}}{\partial s} - \frac{\partial \hat{f}_{ST}^*}{\partial s} \right|(s, t) + \left| \frac{\partial \hat{f}_{ST}}{\partial t} - \frac{\partial \hat{f}_{ST}^*}{\partial t} \right|(s, t)} = O_P(1). \quad (\text{A.10})$$

From standard kernel arguments, it is known that, if $nh^4 \rightarrow \infty$ as $n \rightarrow \infty$, $\partial \hat{f}_{ST}^*(s, t)/\partial s$ is a consistent estimator of $\partial f_{ST}(s, t)/\partial s$, with variance $O((nh^4)^{-1})$ and bias $O(h^2)$. In a way very similar to the proof of Proposition 3.1, one can show that resorting to pseudo-observations does not change the statistical properties of that estimator of $\partial f_{ST}(s, t)/\partial s$. Hence $|\partial \hat{f}_{ST}^*(s, t)/\partial s - \partial \hat{f}_{ST}(s, t)/\partial s| = o_P((nh^4)^{-1/2})$ (and same for the partial derivatives with respect to t). It follows that $\left| \hat{\psi}_{ST}(s, t) - \hat{\psi}_{ST}^*(s, t) \right| = o_P((nh^4)^{-1/2})$. With the factor h^2 in (A.9), this means that the second term is of order $o_P(n^{-1/2})$, which is obviously $o_P((nh^2)^{-1/2})$. Then,

$$|\tilde{f}_{ST}^{(1)}(s, t) - \tilde{f}_{ST}^{*(1)}(s, t)| = o_P((nh^2)^{-1/2}),$$

and since this order carries over to $|\tilde{c}^{(\tau,1)}(u, v) - \tilde{c}^{*(\tau,1)}(u, v)|$ in a straightforward way, (3.13) also holds with $\tilde{c}^{(\tau,1)}(u, v)$ instead of $\tilde{c}^{*(\tau,1)}(u, v)$. \square

Proof of Theorem 3.3

It is very similar to the proof of Theorem 3.2, based on the analogue of (3.15) for $\tilde{f}_{ST}^{(2)}(s, t)$ (i.e. the bivariate version of equation (5.2) in Hjort and Jones (1996)). It is, therefore, omitted. The reason why a condition on h stronger than previously is needed is that here, the analogue of (A.10) involves the second order partial derivatives of \hat{f}_{ST} . To have those consistent for the corresponding partial derivatives of f_{ST} , indeed, requires $nh^6 \rightarrow \infty$. \square

References

- Autin, F., Le Pennec, E. and Tribouley, K. (2010), Thresholding methods to estimate the copula density, *J. Multivariate Anal.*, 101, 200-222.
- Behnen, K., Huskova, M. and Neuhaus, G. (1985), Rank estimators of scores for testing independence, *Statist. Decisions*, 3, 239-262.
- Blumentritt, T., On Copula Density Estimation and Measures of Multivariate Association, PhD Dissertation, Universität zu Köln, 2011.

- Bouezmarni, T., Rombouts, J.V.K. and Taamouti, A. (2010), Asymptotic properties of the Bernstein density copula estimator for α -mixing data, *J. Multivariate Anal.*, 101, 1-10.
- Bouezmarni, T., El Gouch, A. and Taamouti, A. (2013), Bernstein estimator for unbounded copula densities, *Statistics and Risk Modeling*, 4, 343-360.
- Bücher, A. and Volgushev, S. (2013), Empirical and sequential empirical copula processes under serial dependence, *J. Multivariate Anal.*, 119, 61-70.
- Chacón J.E., Duong, T. and Wand, M.P. (2011), Asymptotics for general multivariate kernel density derivative estimators, *Statist. Sinica*, 21, 807-840.
- Charpentier, A., Fermanian, J.-D. and Scaillet, O. (2007), The estimation of copulas: theory and practice, in: J. Rank (Ed.), *Copulas: From Theory to Application in Finance*, Risk Publications, London, pp. 35-60.
- Chen, S.X. (1999), Beta kernels estimators for density functions, *Comput. Statist. Data Anal.*, 31, 131-145.
- Chen, S.X. and Huang, T.-M. (2007), Nonparametric estimation of copula functions for dependence modelling, *Canad. J. Statist.*, 35, 265-282.
- Chen, X., Fan, J., Pouzo, D. and Ying, Z. (2010), Estimation and model selection of semiparametric multivariate survival functions under general censorship, *J. Econometrics*, 157, 129-142.
- Deheuvels, P. (1979), La fonction de dépendance empirique et ses propriétés, *Acad. Roy. Belg Bull. Cl. Sci.*, 65, 274-292.
- Denuit, M., Purcaru, O. and Van Keilegom, I. (2006), Bivariate Archimedean copula modelling for censored data in non-life insurance, *Journal of Actuarial Practice*, 13, 5-32.
- Devroye, L. and Györfi, L., *Nonparametric Density Estimation: the L1 View*, Wiley, 1985.
- Duong, T. and Hazelton, M.L. (2003), Plug-in bandwidth matrices for bivariate kernel density estimation, *J. Nonparametr. Stat.*, 15, 17-30.
- Duong, T. and Hazelton, M.L. (2005), Convergence rates for unconstrained bandwidth matrix selectors in multivariate kernel density estimation, *J. Multivariate Anal.*, 93, 417-433.
- Einmahl, J. and Ruymgaart, F. (1987), The almost sure behavior of the oscillation modulus of the multivariate empirical process, *Statist. Probab. Lett.*, 6, 87-96.

- Embrechts, P. (2009), Copulas: A personal view, *Journal of Risk and Insurance*, 76, 639-650.
- Fan, J. and Gijbels, I., *Local Polynomial Modelling and Its Applications*, Chapman and Hall/CRC, 1996.
- Faugeras, O. (2009), A quantile-copula approach to conditional density estimation, *J. Multivariate Anal.*, 100, 2083-2099.
- Fermanian, J.-D. and Scaillet, O. (2003), Nonparametric estimation of copulas for time series, *J. Risk*, 5, 25-54.
- Fermanian, J.-D., Radulovic, D. and Wegkamp, M. (2004), Weak convergence of empirical copula processes, *Bernoulli*, 10, 847-860.
- Fermanian, J.-D. (2005), Goodness-of-fit tests for copulas, *J. Multivariate Anal.*, 95, 119-152.
- Frees, E. and Valdez, E. (1998), Understanding relationships using copulas, *N. Am. Actuar. J.*, 2, 1-25.
- Geenens, G. (2014), Probit transformation for kernel density estimation on the unit interval, *J. Amer. Statist. Assoc.*, *in press*.
- Genest, C. and Rémillard, B. (2004), Test of independence and randomness based on the empirical copula process, *Test*, 13, 335-369.
- 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, 347-368.
- Genest, C., Rémillard, B. and Beaudoin, D. (2009a), Goodness-of-fit tests for copulas: A review and a power study, *Insurance: Mathematics and Economics*, 44, 199-213.
- Genest, C., Masiello, E. and Tribouley, K. (2009b), Estimating copula densities through wavelets, *Insurance: Mathematics and Economics*, 44, 170-181.
- Genest, C. and Segers, J. (2010), On the covariance of the asymptotic empirical copula process, *J. Multivariate Anal.*, 101, 1837-1845.
- Gijbels, I. and Mielniczuk, J. (1990), Estimating the density of a copula function, *Comm. Statist. Theory Methods*, 19, 445-464.
- Gijbels, I., Omelka, M. and Sznajder, D. (2010), Positive quadrant dependence tests for copulas, *Canad. J. Statist.*, 38, 555-581.

- Gudendorf, G. and Segers, J. (2012), Nonparametric estimation of multivariate extreme-value copulas, *J. Statist. Plann. Inference*, 142, 3073-3085.
- Hall, P. and Neumeier, N. (2006), Estimating a bivariate density when there are extra data on one or both components, *Biometrika*, 93, 439-450.
- Härdle, W. and Okhrin, O. (2009), De copulis non est disputandum – copulae: an overview, *AStA Adv. Stat. Anal.*, 94, 1-31.
- Hjort, N.L. and Jones, M.C. (1996), Locally parametric nonparametric density estimation, *Ann. Statist.*, 24, 1619-1647.
- Janssen, P., Swanepoel, J. and Veraverbeke, N. (2013), A note on the asymptotic behavior of the Bernstein estimator of the copula density, *J. Multivariate Anal.*, 124, 480-487.
- Joe, H., *Multivariate models and dependence concepts*, Chapman and Hall, London, 1997.
- Jones, M.C. and Signorini, D.F. (1997), A comparison of higher-order bias kernel density estimators, *J. Amer. Statist. Assoc.*, 92, 1063-1073.
- Kauermann, G., Schellhase, C. and Ruppert, D. (2013), Flexible Copula Density Estimation with Penalized Hierarchical B-splines, *Scand. J. Statist.*, 40, 685-705.
- Klugman, S. A. and Parsa, R. (1999), Fitting bivariate loss distributions with copulas. *Insurance: Mathematics and Economics*, 24, 139-148.
- Lawlor, G.R. (2012), A L'Hospital rule for multivariable functions, *Manuscript*, arXiv:1209.0363v1.
- Li, B. and Genton, M. (2013), Nonparametric identification of copula structures, *J. Amer. Statist. Assoc.*, 108, 666-675.
- Loader, C.R. (1996), Local likelihood density estimation, *Ann. Statist.*, 24, 1602-1618.
- Loader, C.R., *Local Regression and Likelihood*, Springer, 1999.
- Lopez-Paz, D., Hernández-Lobato, J.M. and Schölkopf, B. (2013), Semi-supervised domain adaptation with non-parametric copulas, *Manuscript*.
- Mack, Y.P. and Rosenblatt, M. (1979), Multivariate k -nearest neighbor density estimates, *J. Multivariate Anal.*, 9, 1-15.

- Marron, J.S. and Ruppert D. (1994), Transformations to reduce boundary bias in kernel density estimation, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 56, 653-671.
- Nelsen, R.B., *An introduction to copulas*, Springer Verlag, New York, 2006.
- Omelka, M., Gijbels, I. and Veraverbeke, N. (2009), Improved kernel estimation of copulas: weak convergence and goodness-of-fit testing, *Ann. Stat.*, 37, 3023-3058.
- Qu, L. and Yin, W. (2012), Copula density estimation by total variation penalized likelihood with linear equality constraints, *Comput. Statist. Data Anal.*, 56, 384-398.
- Salmon, F. (2009), Recipe for disaster: the formula that killed Wall Street, *Wired Magazine*, February 23 2009.
- Scaillet, O. (2007), Kernel-based goodness-of-fit tests for copulas with fixed smoothing parameters, *J. Multivariate Anal.*, 98, 533-543.
- Segers, J. (2012), Asymptotics of empirical copula processes under non-restrictive smoothness assumptions, *Bernoulli*, 18, 764-782.
- Shen, X., Zhu, Y., Song, L. (2008), Linear B-spline copulas with applications to nonparametric estimation of copulas, *Comput. Statist. Data Anal.*, 52, 3806-3819.
- Sklar, A. (1959), Fonctions de répartition à n dimensions et leurs marges, *Publications de l'Institut de Statistique de l'Université de Paris*, 8, 299-331.
- Tsukuhara, H. (2005), Semiparametric estimation in copula models, *Canad. J. Statist.*, 33, 357-375.
- Wand, M.P. and Jones, M.C., *Kernel Smoothing*, Chapman and Hall, 1995.