



Improving estimation for asymptotically independent bivariate extremes via global estimators for the angular dependence function

C. J. R. Murphy-Barltrop^{1,2} · J. L. Wadsworth³ · E. F. Eastoe³

Received: 23 March 2023 / Revised: 25 June 2024 / Accepted: 13 July 2024
© The Author(s) 2024

Abstract

Modelling the extremal dependence of bivariate variables is important in a wide variety of practical applications, including environmental planning, catastrophe modelling and hydrology. The majority of these approaches are based on the framework of bivariate regular variation, and a wide range of literature is available for estimating the dependence structure in this setting. However, such procedures are only applicable to variables exhibiting asymptotic dependence, even though asymptotic independence is often observed in practice. In this paper, we consider the so-called ‘angular dependence function’; this quantity summarises the extremal dependence structure for asymptotically independent variables. Until recently, only pointwise estimators of the angular dependence function have been available. We introduce a range of global estimators and compare them to another recently introduced technique for global estimation through a systematic simulation study, and a case study on river flow data from the north of England, UK.

Keywords Bivariate extremes · Dependence modelling · Asymptotic independence · Angular dependence function

AMS 2000 Subject Classification 62G32 · 62H12

✉ C. J. R. Murphy-Barltrop
callum.murphy-barltrop@tu-dresden.de

¹ Technische Universität Dresden, Institut Für Mathematische Stochastik, Helmholtzstraße 10, 01069 Dresden, Germany

² Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig, Germany

³ School of Mathematical Sciences, Lancaster University, LA1 4YF Lancaster, UK

1 Introduction

Bivariate extreme value theory is a branch of statistics that deals with the modelling of dependence between the extremes of two variables. This type of analysis is useful in a variety of fields, including finance (Castro-Camilo et al. 2018), engineering (Ross et al. 2020), and environmental science (Brunner et al. 2016), where understanding and predicting the behaviour of rare, high-impact events is important.

In certain applications, interest lies in understanding the risk of observing simultaneous extreme events at multiple locations; for example, in the context of flood risk modelling, widespread flooding can result in damaging consequences to properties, businesses, infrastructure, communications and the economy (Lamb et al. 2010; Keef et al. 2013b). To support resilience planning, it is imperative to identify locations at high risk of joint extremes.

Classical theory for bivariate extremes is based on the framework of regular variation. Given a random vector (X, Y) with standard exponential margins, we say that (X, Y) is bivariate regularly varying if, for any measurable $B \subset [0, 1]$,

$$\lim_{r \rightarrow \infty} \Pr(V \in B, R > sr \mid R > r) = H(B)s^{-1}, \quad s \geq 1, \quad (1.1)$$

with $R := e^X + e^Y$, $V := e^X/R$ and $H(\partial B) = 0$, where ∂B is the boundary of B (Resnick 1987). Note that bivariate regular variation is most naturally expressed on standard Pareto margins, and the mapping $(X, Y) \mapsto (e^X, e^Y)$ performs this transformation. We refer to R and V as radial and angular components, respectively. Equation 1.1 implies that for the largest radial values, the radial and angular components are independent. Furthermore, the quantity H , which is known as the spectral measure, must satisfy the moment constraint $\int_0^1 v dH(v) = 1/2$.

The spectral measure summarises the extremal dependence of (X, Y) , and a wide range of approaches exist for its estimation (e.g., Einmahl and Segers 2009; de Carvalho and Davison 2014; Eastoe et al. 2014). Equivalently, one can consider Pickands' dependence function (Pickands 1981), which has a direct relationship to H via

$$A(t) = \int_0^1 \max\{vt, (1-v)(1-t)\} 2dH(v), \quad t \in [0, 1],$$

where A is a convex function satisfying $\max(t, 1-t) \leq A(t) \leq 1$. This function again captures the extremal dependence of (X, Y) , and many approaches also exist for its estimation (e.g., Guillothe and Perron 2016; Marcon et al. 2016; Vettori et al. 2018). Moreover, estimation procedures for the spectral measure and Pickands' dependence function encompass a wide range of statistical methodologies, with parametric, semi-parametric, and non-parametric modelling techniques proposed in both Bayesian and frequentist settings.

However, methods based on bivariate regular variation are limited in the forms of extremal dependence they can capture. This dependence can be classified through the coefficient χ (Joe 1997), defined as

$$\chi := \lim_{u \rightarrow \infty} \Pr(Y > u \mid X > u) \in [0, 1],$$

where this limit exists. If $\chi > 0$, then X and Y are asymptotically dependent, and the most extreme values of either variable can occur simultaneously. If $\chi = 0$, X and Y are asymptotically independent, and the most extreme values of either variable occur separately.

Under asymptotic independence, the spectral measure H places all mass on the points $\{0\}$ and $\{1\}$; equivalently, $A(t) = 1$ for all $t \in [0, 1]$. Consequently, for this form of dependence, the framework given in Eq. 1.1 is degenerate and is unable to accurately extrapolate into the joint tail (Ledford and Tawn 1996, 1997). Practically, an incorrect assumption of asymptotic dependence between two variables is likely to result in an overly conservative estimate of joint risk.

To overcome this limitation, several models have been proposed that can capture both classes of extremal dependence. The first was given by Ledford and Tawn (1996), in which they assume that as $u \rightarrow \infty$, the joint tail can be represented as

$$\Pr(X > u, Y > u) = \Pr\{\min(X, Y) > u\} = L(e^u)e^{-u/\eta}, \tag{1.2}$$

where L is a slowly varying function at infinity, i.e., $\lim_{u \rightarrow \infty} L(cu)/L(u) = 1$ for $c > 0$, and $\eta \in (0, 1]$. The quantity η is termed the coefficient of tail dependence, with $\eta = 1$ and $\lim_{u \rightarrow \infty} L(u) > 0$ corresponding to asymptotic dependence and either $\eta < 1$ or $\eta = 1$ and $\lim_{u \rightarrow \infty} L(u) = 0$ corresponding to asymptotic independence. Many extensions to this approach exist (e.g., Ledford and Tawn 1997; Resnick 2002; Ramos and Ledford 2009); however, all such approaches are only applicable in regions where both variables are large, limiting their use in many practical settings. Since many extremal bivariate risk measures, such as environmental contours (Haselsteiner et al. 2021) and return curves (Murphy-Bartrop et al. 2023), are defined both in regions where both variables are extreme and in regions where only one variable is extreme, methods based on Eq. 1.2 are inadequate for their estimation.

Several copula-based models have been proposed that can capture both classes of extremal dependence, such as those given in Coles and Pauli (2002), Wadsworth et al. (2017) and Huser and Wadsworth (2019). Unlike Eq. 1.2, these can be used to evaluate joint tail behaviour in all regions where at least one variable is extreme. However, these techniques typically require strong assumptions about the parametric form of the bivariate distribution, thereby offering reduced flexibility.

Heffernan and Tawn (2004) proposed a modelling approach, known as the conditional extremes model, which also overcomes the limitations of the framework described in Eq. 1.2. This approach assumes the existence of normalising functions $a : \mathbb{R}_+ \rightarrow \mathbb{R}$ and $b : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$\lim_{u \rightarrow \infty} \Pr \left[\{Y - a(X)\}/b(X) \leq z, X - u > x \mid X > u \right] = D(z)e^{-x}, \quad x > 0, \tag{1.3}$$

where D is a non-degenerate distribution function that places no mass at infinity. Note that the choice of conditioning on $X > u$ is arbitrary, and an equivalent formulation exists for normalised X given $Y > u$. This framework can capture both asymptotic dependence and asymptotic independence, with the former arising when $a(x) = x$ and $b(x) = 1$, and can also be used to describe extremal behaviour in regions where only one variable is large.

Finally, Wadsworth and Tawn (2013) proposed a general extension of Eq. 1.2. As $u \rightarrow \infty$, they assume that for any $(\beta, \gamma) \in \mathbb{R}_+^2 \setminus \{\mathbf{0}\}$,

$$\Pr(X > \beta u, Y > \gamma u) = L(e^\mu; \beta, \gamma) e^{-\kappa(\beta, \gamma)u}, \tag{1.4}$$

where $L(\cdot; \gamma, \beta)$ is slowly varying and the function κ provides information about the joint tail behaviour of (X, Y) . One can observe that Eq. 1.2 is a special case of Eq. 1.4 with $\beta = \gamma$. The dependence function κ satisfies several theoretical properties: for instance, it is non-decreasing in each argument, satisfies the lower bound $\kappa(\beta, \gamma) \geq \max\{\beta, \gamma\}$, and is homogeneous of order 1, i.e., $\kappa(h\beta, h\gamma) = h\kappa(\beta, \gamma)$ for any $h > 0$. Setting $w := \beta/(\beta + \gamma) \in [0, 1]$, the latter property implies that $\kappa(\beta, \gamma) = (\beta + \gamma)\kappa(w, 1 - w)$, motivating the definition of the so-called angular dependence function (ADF) $\lambda(w) = \kappa(w, 1 - w)$, $w \in [0, 1]$. Using this representation, Eq. 1.4 can be rewritten as

$$\Pr(\min\{X/w, Y/(1 - w)\} > u) = L(e^\mu; w) e^{-\lambda(w)u}, \quad w \in [0, 1], \quad \lambda(w) \geq \max(w, 1 - w), \tag{1.5}$$

as $u \rightarrow \infty$, where $L(\cdot; w)$ is slowly varying. The ADF generalises the coefficient η , with $\eta = 1/\{2\lambda(0.5)\}$. This extension captures both extremal dependence regimes, with asymptotic dependence implying the lower bound, i.e., $\lambda(w) = \max(w, 1 - w)$ for all $w \in [0, 1]$. Evaluation of the ADF for rays w close to 0 and 1 corresponds to regions where one variable is larger than the other.

The ADF can be viewed as the counterpart of the Pickands' dependence function for asymptotically independent variables, and shares many of its theoretical properties (Wadsworth and Tawn 2013). Specifically, $\lambda(0) = \lambda(1) = 1$ and $\max(w, 1 - w) \leq \lambda(w)$, although there is no requirement for $\lambda(w)$ to be convex, or that it is bounded above, unlike $A(t)$. There do, however, exist shape constraints that λ must satisfy; see Sect. 3 for further details. The ADF can be used to differentiate between different forms of asymptotic independence, with both positive and negative associations captured, alongside complete independence, which implies $\lambda(w) = 1$ for all $w \in [0, 1]$. Figure 1 illustrates the ADFs for three copulas. We observe a variety in shapes, corresponding




Fig. 1 The true ADFs (given in red) for three example copulas. Left: bivariate Gaussian copula with coefficient $\rho = 0.5$. Centre: inverted logistic copula with dependence parameter $r = 0.8$. Right: logistic copula with dependence parameter $r = 0.8$. The lower bound for the ADF is denoted by the black dotted line

to differing degrees of positive extremal dependence in the underlying copulas. The weakest dependence is observed for the inverted logistic copula, while the ADF for the asymptotically dependent logistic copula is equal to the lower bound.

Despite these modelling advances, the majority of approaches for quantifying the risk of bivariate extreme events still require bivariate regular variation. Many of the procedures that do allow for asymptotic independence use the conditional extremes model of Eq. 1.3 despite some well known limitations of this approach (Liu and Tawn 2014).

One particular application of the model described in Eq. 1.5 is the estimation of so-called bivariate return curves, $RC(p) := \{(x, y) \in \mathbb{R}^2 : \Pr(X > x, Y > y) = p\}$, which requires knowledge of extremal dependence in regions where either variable is large; see Sect. 6.4. Murphy-Barltrop et al. (2023) obtain estimates of return curves, finding that estimates derived using Eq. 1.5 were preferable to those from the conditional extremes model. Mhalla et al. (2019b) and Murphy-Barltrop and Wadsworth (2024) also provide non-stationary extensions and inference methods for the ADF.

In this paper, we propose a global methodology for ADF estimation in order to improve extrapolation into the joint upper tail for bivariate random vectors exhibiting asymptotic independence. Until recently, the ADF has been estimated only in a pointwise manner using the Hill estimator (Hill 1975) on the tail of $\min\{X/w, Y/(1-w)\}$, resulting in unrealistic rough functional estimates and, as we demonstrate in Sect. 5, high degrees of variability. Further, Murphy-Barltrop et al. (2023) showed that pointwise ADF estimates result in non-smooth return curve estimates, which are again unrealistic.

The first smooth ADF estimator was proposed recently in Simpson and Tawn (2022) based on a theoretical link between a limit set derived from the shape of appropriately scaled sample clouds and the ADF (Nolde and Wadsworth 2022). The authors introduce global estimation techniques for the limit set, from which smooth ADF estimates follow; see Sect. 2 for further details.

We introduce several novel smooth ADF estimators, and compare their performance with the pointwise Hill estimator, as well as the estimator given in Simpson and Tawn (2022). In Sect. 2, we review the literature on ADF estimation. In Sect. 3, we introduce new theoretical results that the ADF must satisfy to be valid. In Sect. 4, we introduce a range of novel estimators, and select tuning parameters for each proposed estimation technique. In Sect. 5, we compare each of the available estimators through a systematic simulation study, finding certain estimators to be favourable over others. A subset of estimators are then applied to river flow data sets in Sect. 6 and used to obtain estimates of return curves for different combinations of river gauges. We conclude in Sect. 7 with a discussion.

2 Existing techniques for ADF estimation

In this section, we introduce existing estimators for the ADF, with (X, Y) denoting a random vector with standard exponential margins throughout. To begin, for any ray $w \in [0, 1]$, define the min-projection at w as $T_w := \min\{X/w, Y/(1-w)\}$. Equation 1.5 implies that for any $w \in [0, 1]$ and $t > 0$,

$$\Pr(T_w > u + t \mid T_w > u) = \frac{L(e^{u+t}; w)}{L(e^u; w)} e^{-\lambda(w)t} \rightarrow e^{-\lambda(w)t} = t_*^{-\lambda(w)}, \quad (2.1)$$

as $u \rightarrow \infty$, with $t_* := e^t$. Since the expression in Eq. 2.1 has a univariate regularly varying tail with positive index, Wadsworth and Tawn (2013) propose using the Hill estimator (Hill 1975) to obtain a pointwise estimator of the ADF; we denote this ‘base’ estimator $\hat{\lambda}_H$. A major drawback of this technique is that the estimator is pointwise, that is, $\lambda(w)$ is estimated separately for each w , leading to rough and often unrealistic estimates of the ADF. In particular, no information is shared across different rays, increasing the variability in the resulting estimates. Furthermore, this estimator need not satisfy the theoretical constraints on the ADF identified in Wadsworth and Tawn (2013), such as the endpoint conditions $\lambda(0) = \lambda(1) = 1$.

Simpson and Tawn (2022) recently proposed a novel estimator for the ADF using a theoretical link with the limiting shape of scaled sample clouds. Let $C_n := \{(X_i, Y_i) / \log n; i = 1, \dots, n\}$ denote n scaled, independent copies of (X, Y) . Nolde and Wadsworth (2022) explain how, as $n \rightarrow \infty$, the asymptotic shape of C_n provides information on the underlying extremal dependence structure. In many situations, C_n converges onto the compact limit set $G = \{(x, y) : g(x, y) \leq 1\} \subseteq [0, 1]^2$, where g is the gauge function of G . A sufficient condition for this convergence to occur is that the joint density, f , of (X, Y) exists, and that

$$-\log f(tx, ty) \sim tg(x, y), \quad x, y \geq 0, \quad t \rightarrow \infty, \quad (2.2)$$

for continuous g . Following Nolde (2014), we also define the unit-level, boundary set $\partial G = \{(x, y) : g(x, y) = 1\} \subset [0, 1]^2$. Given fixed margins, the shapes of G , and hence ∂G , are completely determined by the extremal dependence structure of (X, Y) . Furthermore, Nolde and Wadsworth (2022) show that the shape of G is also directly linked to the modelling frameworks described in Eqs. 1.2, 1.3 and 1.5, as well as the approach of Simpson et al. (2020). In particular, letting $R_w := (w / \max(w, 1 - w), \infty) \times ((1 - w) / \max(w, 1 - w), \infty)$ for all $w \in [0, 1]$, we have that

$$\lambda(w) = \max(w, 1 - w) \times r_w^{-1}, \quad (2.3)$$

where

$$r_w = \min\{r \in [0, 1] : rR_w \cap G = \emptyset\}.$$

The boundary sets ∂G for each of the copulas in Fig. 1 are given in Fig. 2, alongside the coordinates $(w/\lambda(w), (1-w)/\lambda(w))$ for all $w \in [0, 1]$; these coordinates represent the relationship between G and the ADF via Eq. 2.3. One can again observe the variety in shapes. For the asymptotically dependent logistic copula, we have that $(1, 1) \in \partial G$; this is true for all asymptotically dependent bivariate random vectors with a limit set.

In practice, both the limit set, G , and its boundary, ∂G , are unknown. Simpson and Tawn (2022) propose an estimator for ∂G , which is then used to derive an estimator $\hat{\lambda}_{ST}$ for the ADF via Eq. 2.3. The resulting estimator $\hat{\lambda}_{ST}$ was shown to outperform $\hat{\lambda}_H$ in a wide range of scenarios (Simpson and Tawn 2022).




Fig. 2 The boundary set ∂G (given in red) for three example copulas, with coordinate limits denoted by the black dotted lines and the blue lines representing the coordinates $(w/\lambda(w), (1 - w)/\lambda(w))$ for all $w \in [0, 1]$. Left: bivariate Gaussian copula with coefficient $\rho = 0.5$. Centre: inverted logistic copula with dependence parameter $r = 0.8$. Right: logistic copula with dependence parameter $r = 0.8$

Estimation of ∂G uses an alternative radial-angular decomposition of (X, Y) , with $R^* := X + Y$ and $V^* := X/(X + Y)$. Simpson and Tawn (2022) assume the tail of $R^* | V^* = v^*$, $v^* \in [0, 1]$, follows a generalised Pareto distribution (Davison and Smith 1990) and then use generalised additive models to capture trends over angles in both the threshold and generalised Pareto distribution scale parameter (Youngman 2019). Next, high quantile estimates from the conditional distributions $R^* | V^* = v^*$, $v^* \in [0, 1]$ are computed using the fitted generalised Pareto distributions. They are then transformed back to the original scale using $X = R^*V^*$ and $Y = R^*(1 - V^*)$ and finally scaled onto the set $[0, 1]^2$ to give an estimate of ∂G ; see Simpson and Tawn (2022) for further details.

Wadsworth and Campbell (2024) also provide methodology for estimation of ∂G , though their focus is on estimation of tail probabilities more generally, including in dimensions greater than two. Furthermore, their approach requires prior selection of a parametric form for g . We therefore restrict our attention to the work of Simpson and Tawn (2022) as their main focus is semi-parametric estimation for ∂G in two dimensions.

When applying the estimators $\hat{\lambda}_H$ and $\hat{\lambda}_{ST}$ in Sect. 5 and 6, we use the tuning parameters suggested in the original approaches. In the case of $\hat{\lambda}_H$, we set u to be the empirical 90% quantile of T_w . The default tuning parameters for $\hat{\lambda}_{ST}$ can be found in Simpson and Tawn (2022), and example estimates of the set ∂G obtained using the suggested parameters are given in the Supplementary Material. For calculating this estimator, we used the code available at <https://github.com/essimpson/self-consistent-inference>.

3 Novel theoretical results pertaining to the ADF

In this section, we outline some new theoretical results on the ADF. These results impose shape constraints that the ADF must satisfy, and follow from the properties of κ introduced in Sect. 1.

Proposition 3.1 For any $w_1, w_2 \in [0, 1]$ such that $w_1 \leq w_2$, we have

$$w_1/\lambda(w_1) \leq w_2/\lambda(w_2) \quad \text{and} \quad (1 - w_1)/\lambda(w_1) \geq (1 - w_2)/\lambda(w_2).$$

Proof For the first statement, the proof is trivial if at least one of w_1 or w_2 equals 0. Therefore, without loss of generality, assume that $w_1, w_2 > 0$. Recall $\lambda(w) = \kappa(w, 1 - w)$, and that κ is homogeneous of order 1 and monotonic in both arguments. Setting $t := w_1/w_2 \in (0, 1]$, we have

$$\begin{aligned} \frac{w_1}{w_2} \lambda(w_2) &= t\kappa(w_2, 1 - w_2) \\ &= \kappa(w_1, t(1 - w_2)) \\ &\leq \kappa(w_1, t(1 - w_1)) \\ &\leq \kappa(w_1, 1 - w_1) \\ &= \lambda(w_1), \end{aligned}$$

implying $w_1/\lambda(w_1) \leq w_2/\lambda(w_2)$. For the second argument, the proof is again trivial if $w_2 = 1$ or $w_1 = w_2 = 1$, so assume $w_1, w_2 < 1$. Setting $t := (1 - w_2)/(1 - w_1) \in (0, 1]$, similar reasoning shows that $(1 - w_1)/\lambda(w_1) \geq (1 - w_2)/\lambda(w_2)$, completing the proof. \square

Remark The intuition behind Proposition 3.1 comes from considering the relationship between the ADF and the boundary set ∂G described in Eq. 2.3. Considering the sets described by the blue lines in Fig. 2, this implies that the x -coordinates (y -coordinates) of the sets must be increasing (decreasing) as the ray w increases, as is clear from the figure.

Proposition 3.1 also implies several interesting properties that the ADF must satisfy.

Corollary 3.1 Suppose there exists $w^* \in [0, 0.5]$, or $w^* \in [0.5, 1]$, such that $\lambda(w^*) = \max(w^*, 1 - w^*)$. Then $\lambda(w) = \max(w, 1 - w)$ for all $w \in [0, w^*]$, or $w \in [w^*, 1]$.

Proof Considering the first case $w^* \in [0, 0.5]$, Proposition 3.1 gives that $(1 - w)/\lambda(w) \geq (1 - w^*)/\lambda(w^*) = (1 - w^*)/(1 - w^*) = 1$ for any $w \in [0, w^*]$, implying $\max(w, 1 - w) = (1 - w) \geq \lambda(w)$. Since $\lambda(w) \geq \max(w, 1 - w)$, we must therefore have $\lambda(w) = \max(w, 1 - w)$. The same reasoning applies for $w \in [w^*, 1]$ when $w^* \geq 0.5$ and $\lambda(w^*) = \max(w^*, 1 - w^*)$. \square

Corollary 3.1 states that if the ADF equals the lower bound for any angle in the interval $[0, 0.5]$ (or the interval $[0.5, 1]$), then it must also equal the lower bound for all angles less (greater) than this angle. This has further implications when we consider the conditional extremes modelling framework described in Eq. 1.3. Let $a_{y|x}$ and $a_{x|y}$ be the normalising functions for conditioning on the events $X > u$ and $Y > u$ respectively, and let $\alpha_{y|x} := \lim_{u \rightarrow \infty} a_{y|x}(u)/u$ and $\alpha_{x|y} := \lim_{u \rightarrow \infty} a_{x|y}(u)/u$, with

$\alpha_{y|x}, \alpha_{x|y} \in [0, 1]$. From Nolde and Wadsworth (2022), we have that $g(1, \alpha_{y|x}) = 1$ and $g(\alpha_{x|y}, 1) = 1$, with g defined as in Eq. 2.2, and $\alpha_{y|x}, \alpha_{x|y}$ are the maximum such values satisfying these equations. Assuming that the values of $\alpha_{y|x}$ and $\alpha_{x|y}$ are known, we have the following result.

Corollary 3.2 For all $w \in [0, \alpha_{x|y}^*] \cup [\alpha_{y|x}^*, 1]$, with $\alpha_{x|y}^* := \alpha_{x|y}/(1 + \alpha_{x|y})$ and $\alpha_{y|x}^* := 1/(1 + \alpha_{y|x})$, we have $\lambda(w) = \max(w, 1 - w)$.

Proof To begin, consider the ray $\alpha_{x|y}^* \in [0, 0.5]$ and observe that $(\alpha_{x|y}, 1) \in \partial G$. From this, one can see that $(\alpha_{x|y}, 1) \in R_{\alpha_{x|y}^*}$. Equation 2.3 therefore implies that $r_{\alpha_{x|y}^*} = 1$, and hence $\lambda(\alpha_{x|y}^*) = \max(\alpha_{x|y}^*, 1 - \alpha_{x|y}^*)$. From Corollary 3.1, it follows that $\lambda(w) = \max(w, 1 - w)$ for all $w \in [0, \alpha_{x|y}^*]$. Considering the ray $\alpha_{y|x}^* \in [0.5, 1]$ in an analogous manner, the result follows. \square

Corollaries 3.1 and 3.2 are illustrated in Fig. 3 for a Gaussian copula with $\rho = 0.5$. Here, $\alpha_{x|y} = \alpha_{y|x} = 0.25$, implying $\lambda(w) = \max(w, 1 - w)$ for all $w \in [0, 0.2] \cup [0.8, 1]$; these rays correspond to the blue lines in the figure. One can observe that for any region R_w defined along either of the blue lines (such as the shaded regions illustrated for $w = 0.1$ and $w = 0.9$), we have that $r_w = 1$, since these regions will intersect ∂G at either the coordinates $(0.25, 1)$ or $(1, 0.25)$.

Finally, Proposition 3.1 also implies constraints on the derivative of the ADF, where this exists.

Corollary 3.3 Let $\lambda'(w)$ denote the derivative of the ADF at w , where it exists. For all $w \in (0, 1)$, we have that

$$-\lambda(w)/(1 - w) \leq \lambda'(w) \leq \lambda(w)/w.$$

Proof Given any $w \in (0, 1)$ and $\epsilon > 0$ such that $\epsilon \leq \min\{w, 1 - w\}$, Proposition 3.1 implies that $w/\lambda(w) \leq (w + \epsilon)/\lambda(w + \epsilon)$ and $(w - \epsilon)/\lambda(w - \epsilon) \leq w/\lambda(w)$, which can be rearranged to give

$$\frac{\lambda(w + \epsilon) - \lambda(w)}{\epsilon} \leq \lambda(w)/w, \quad \frac{\lambda(w) - \lambda(w - \epsilon)}{\epsilon} \leq \lambda(w)/w.$$

Taking the limits as $\epsilon \rightarrow 0^+$, we obtain $\lambda'(w) \leq \lambda(w)/w$. By rearranging the equations $(1 - w)/\lambda(w) \geq (1 - (w + \epsilon))/\lambda(w + \epsilon)$ and $(1 - (w - \epsilon))/\lambda(w - \epsilon) \geq (1 - w)/\lambda(w)$ and taking limits in a similar manner, the result follows. \square

Remark Note that Corollary 3.3 corresponds to the same derivative constraints originally derived in Wadsworth and Tawn (2013), albeit with a different proof.

The results introduced in this section provide necessary conditions that the ADF must satisfy. Furthermore, each of the introduced corollaries follow directly from Proposition 3.1; therefore, if an estimator satisfies the conditions of this proposition, it will automatically satisfy the remaining conditions. We therefore incorporate




Fig. 3 Pictorial illustration of the results described in Corollaries 3.1 and 3.2. The boundary set ∂G , given in red, is from the bivariate Gaussian copula with $\rho = 0.5$, with the points $(1, \alpha_{y|x})$ and $(\alpha_{x|y}, 1)$ given in green. The blue lines represent the rays $w \in [0, \alpha_{x|y}] \cup [\alpha_{y|x}, 1]$, while the yellow and pink shaded regions represent the set R_w for $w = 0.1$ and $w = 0.9$, respectively

Proposition 3.1 into our estimation framework for λ ; see Sect. 4 for further details. We remark that only one of the existing approaches for estimating the ADF (Murphy-Barltrop and Wadsworth 2024) has considered such constraints, although they will automatically be satisfied by ADF estimators derived from valid limit set estimates.

4 Novel estimators for the ADF

Motivated by the goal of global estimation, we propose a range of novel estimators for the ADF. We recall that the ADF and Pickands' dependence function exhibit several theoretical similarities, as listed in Sect. 1, and arise as exponential rate parameters for suitable constructions of structure variables (Mhalla et al. 2019a). We therefore begin by reviewing estimation of the Pickands' dependence function.

Because smooth functional estimation for the ADF is desirable, we restrict our review to approaches for the Pickands' dependence function which achieve this:

notably spline-based techniques (Hall and Tajvidi 2000; Cormier et al. 2014) and techniques that utilise the family of Bernstein-Bézier polynomials (Guillotte and Perron 2016; Marcon et al. 2016, 2017). In this paper, we focus on to the latter category, since spline-based techniques typically result in more complex formulations and a larger number of tuning parameters. Moreover, approaches based on Bernstein-Bézier polynomials have been shown to improve estimator performance across a wide range of copula examples (Vettori et al. 2018). For estimation of Pickands' dependence function, the following family of functions is considered

$$\mathcal{B}_k = \left\{ \sum_{i=0}^k \beta_i \binom{k}{i} w^i (1-w)^{k-i} \mid \boldsymbol{\beta} \in [0, 1]^{k+1}, w \in [0, 1] \right\},$$

where $k \in \mathbb{N}$ denotes the polynomial degree. Note that \mathcal{B}_k is a sub-family from the class of Bernstein-Bézier polynomials. Many approaches assume that the Pickands' dependence function $A \in \mathcal{B}_k$ and propose techniques for estimating the coefficient vector $\boldsymbol{\beta}$, resulting in an estimator $\hat{\boldsymbol{\beta}}$. This automatically ensures $A(t) \leq 1$ for all $t \in [0, 1]$, thereby satisfying the theoretical upper bound of the Pickands' dependence function.

We make a similar assumption about the ADF, and use this to propose novel estimators. However, unlike the Pickands' dependence function, the ADF is unbounded from above, meaning functions in \mathcal{B}_k cannot represent all forms of extremal dependence captured by Eq. 1.5. Moreover, the endpoint conditions $\lambda(0) = \lambda(1) = 1$ are not necessarily satisfied by functions in \mathcal{B}_k . We therefore propose an alternative family of polynomials: given $k \in \mathbb{N}$, let

$$\mathcal{B}_k^* = \left\{ (1-w)^k + \sum_{i=1}^{k-1} \beta_i \binom{k}{i} w^i (1-w)^{k-i} + w^k =: f(w) \mid w \in [0, 1], \boldsymbol{\beta} \in [0, \infty)^{k-1} \right\}. \tag{4.1}$$

Functions in this family are unbounded from above, and $f(0) = f(1) = 1$ for all $f \in \mathcal{B}_k^*$. Note that functions in \mathcal{B}_k^* are not required to satisfy the lower bound of the ADF; this bound is instead imposed in a post-processing procedure, as detailed in Sect. 4.4.

For the remainder of this section, let $\lambda(\cdot; \boldsymbol{\beta}) \in \mathcal{B}_k^*$ represent a form of the ADF from \mathcal{B}_k^* . Interest now lies in estimating the coefficient vector $\boldsymbol{\beta}$, which requires choice of the degree $k \in \mathbb{N}$. This is a trade-off between flexibility and computational complexity; polynomials with small values of k may not be flexible enough to capture all extremal dependence structures, resulting in bias, while high values of k will increase computational burden and parameter variance.

4.1 Composite likelihood approach

One consequence of Eq. 2.1 is that, for all $w \in [0, 1]$, the conditional variable $T_w^* := (T_w - u_w \mid T_w > u_w) \sim \text{Exp}(\lambda(w))$, approximately, for large u_w . The density of this variable is $f_{T_w^*}(t_w^*) \approx \lambda(w)e^{-\lambda(w)t_w^*}$, $t_w^* > 0$, resulting in a likelihood function for

min-projection exceedances of u_w . Let $(\mathbf{x}, \mathbf{y}) := \{(x_i, y_i) : i = 1, \dots, n\}$ denote n independent observations from the joint distribution of (X, Y) . For each $w \in \mathscr{W}$, where \mathscr{W} denotes some finite subset spanning the interval $[0, 1]$, let $\mathbf{t}_w := \{\min(x_i/w, y_i/(1-w)) : i = 1, \dots, n\}$ and take u_w to be the empirical q quantile of \mathbf{t}_w , with q close to 1 and fixed across w . Letting $\mathbf{t}_w^* := \{t_w - u_w \mid t_w \in \mathbf{t}_w, t_w > u_w\}$, we have a set of realisations from the conditional variable T_w^* .

One approach to obtain an estimate of $\lambda(w)$ while considering all $w \in \mathscr{W}$ simultaneously is to use a composite likelihood, in which multiple likelihood components are treated as independent whether or not they are independent. Provided each component is a valid density function, the resulting likelihood function provides unbiased parameter estimates under the true model; see Varin et al. (2011) for further details. For this model, the likelihood function is

$$\mathcal{L}_C(\boldsymbol{\beta}) = \prod_{w \in \mathscr{W}} \prod_{t_w^* \in \mathbf{t}_w^*} \lambda(w; \boldsymbol{\beta}) e^{-\lambda(w; \boldsymbol{\beta}) t_w^*} = \left[\prod_{w \in \mathscr{W}} \lambda(w; \boldsymbol{\beta})^{|\mathbf{t}_w^*|} \right] \times e^{-\sum_{w \in \mathscr{W}} \sum_{t_w^* \in \mathbf{t}_w^*} \lambda(w; \boldsymbol{\beta}) t_w^*}, \quad (4.2)$$

where $|\mathbf{t}_w^*|$ denotes the cardinality of the set \mathbf{t}_w^* . This composite likelihood function has equal weights across all $w \in \mathscr{W}$ (the ‘components’). An estimator of the ADF, $\hat{\lambda}_{CL}$, is $\lambda(\cdot; \hat{\boldsymbol{\beta}}_{CL})$, where $\hat{\boldsymbol{\beta}}_{CL} := \arg \max_{\boldsymbol{\beta} \in [0, \infty)^{k-1}} \mathcal{L}_C(\boldsymbol{\beta})$.

To apply this method in practice, one must first specify a set \mathscr{W} and a probability q . Given some large, odd-valued $m \in \mathbb{N}$, we let $\mathscr{W} := \{i/(m-1) : i = 0, 1, \dots, m-1\}$; this corresponds to a set of equally spaced rays \mathscr{W} spanning the interval $[0, 1]$, with $\{0, 0.5, 1\} \subset \mathscr{W}$. Selection of m and q are discussed in Sect. 4.5. The former is akin to selecting the degree of smoothing, while the latter is analogous to selecting a threshold for the generalised Pareto distribution defined in Sect. 2 in the univariate setting.

4.2 Probability ratio approach

With \mathscr{W} and \mathbf{t}_w defined as in Sect. 4.1, consider two probabilities $q < p < 1$, both close to one. Given any $w \in \mathscr{W}$, let u_w and v_w denote the q and p empirical quantiles of \mathbf{t}_w , respectively. Assuming the distribution function of T_w is strictly monotonic, Eq. 2.1 implies that

$$\frac{1-p}{1-q} = \Pr(T_w > v_w \mid T_w > u_w) \approx e^{-\lambda(w)(v_w - u_w)} \Rightarrow \left| \frac{1-p}{1-q} - e^{-\lambda(w)(v_w - u_w)} \right| \approx 0, \quad (4.3)$$

Similarly to Murphy-Barttrop and Wadsworth (2024), we exploit Eq. 4.3 to obtain an estimator for the ADF. Firstly, we observe that this representation holds for all $w \in \mathscr{W}$, hence

$$\sum_{w \in \mathscr{W}} \left| \frac{1-p}{1-q} - e^{-\lambda(w)(v_w - u_w)} \right| \approx 0.$$

To ensure Eq. 4.3 holds requires careful selection of q and p . This selection also represents a bias-variance trade off: probabilities too small (big) will induce bias (high variability). Moreover, owing to the different rates of convergence to the limiting ADF, a single pair (q, p) is unlikely to be appropriate across all data structures. We instead consider a range of probability pairs simultaneously. Specifically, letting $\{(q_j, p_j) \mid q_j < p_j < 1, 1 \leq j \leq h\}$, $h \in \mathbb{N}$, be pairs of probabilities near one, consider the expression

$$S(\boldsymbol{\beta}) := \sum_{w \in \mathcal{W}} \sum_{j=1}^h \left| \left[\frac{1-p_j}{1-q_j} \right] - e^{-\lambda(w; \boldsymbol{\beta})(v_{w,j}-u_{w,j})} \right|,$$

in which $u_{w,j}$ and $v_{w,j}$ denote q_j and p_j empirical quantiles of \mathbf{t}_w , respectively, for each $j = 1, \dots, h$. Since minimising $S(\boldsymbol{\beta})$ should provide a reasonable estimate of λ over all \mathcal{W} , we set $\hat{\boldsymbol{\beta}}_{PR} = \arg \min_{\boldsymbol{\beta} \in [0, \infty)^{k-1}} S(\boldsymbol{\beta})$ and denote by $\hat{\lambda}_{PR}$ the estimator $\lambda(\cdot; \hat{\boldsymbol{\beta}}_{PR})$. Similarly to $\hat{\lambda}_{CL}$, one must select the sets \mathcal{W} and $\{(q_j, p_j) \mid q_j < p_j < 1, 1 \leq j \leq h\}$, $h \in \mathbb{N}$ prior to applying this estimator; see Sect. 4.5.

4.3 Incorporating knowledge of conditional extremes parameters

Assuming we know the conditional extremes parameters $\alpha_{y|x}, \alpha_{x|y}$, Corollary 3.2 implies that for all $w \in [0, \alpha_{x|y}^*] \cup [\alpha_{y|x}^*, 1]$, with $\alpha_{x|y}^* = \alpha_{x|y}/(1 + \alpha_{x|y})$ and $\alpha_{y|x}^* = 1/(1 + \alpha_{y|x})$, $\lambda(w) = \max(w, 1 - w)$. In this section, we exploit this result to improve estimation of the ADF.

In practice, $\alpha_{x|y}^*$ and $\alpha_{y|x}^*$ are unknown; however, estimates $\hat{\alpha}_{y|x}$ and $\hat{\alpha}_{x|y}$ are commonly obtained using a likelihood function based on a misspecified model for the distribution D in Eq. 1.3 (e.g., Jonathan et al. 2014). The resulting estimates, denoted $\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*$, can be used to approximate the ADF for $w \in [0, \hat{\alpha}_{x|y}^*] \cup [\hat{\alpha}_{y|x}^*, 1]$. What now remains is to combine this with an estimator for $\lambda(w)$ on $[\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*]$.

A crude way to obtain an estimator via this framework would be to set $\lambda(w) = \max(w, 1 - w)$ for $w \in [0, \hat{\alpha}_{x|y}^*] \cup [\hat{\alpha}_{y|x}^*, 1]$ and $\lambda(w) = \hat{\lambda}_H(w)$, $\hat{\lambda}_{CL}(w)$ or $\hat{\lambda}_{PR}(w)$ for $w \in [\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*]$. However, this results in discontinuities at $\hat{\alpha}_{x|y}^*$ and $\hat{\alpha}_{y|x}^*$. Instead, for the smooth estimators, we rescale \mathcal{B}_k^* such that the resulting ADF estimate is continuous for all $w \in [0, 1]$. Consider the set of polynomials

$$\begin{aligned} \tilde{\mathcal{B}}_k = & \left\{ (1 - \hat{\alpha}_{x|y}^*) \left(1 - \frac{v - \hat{\alpha}_{x|y}^*}{\hat{\alpha}_{y|x}^* - \hat{\alpha}_{x|y}^*} \right)^k + \sum_{i=1}^{k-1} \beta_i \binom{k}{i} \left(\frac{v - \hat{\alpha}_{x|y}^*}{\hat{\alpha}_{y|x}^* - \hat{\alpha}_{x|y}^*} \right)^i \left(1 - \frac{v - \hat{\alpha}_{x|y}^*}{\hat{\alpha}_{y|x}^* - \hat{\alpha}_{x|y}^*} \right)^{k-i} + \right. \\ & \left. \hat{\alpha}_{y|x}^* \left(\frac{v - \hat{\alpha}_{x|y}^*}{\hat{\alpha}_{y|x}^* - \hat{\alpha}_{x|y}^*} \right)^k =: f(v) \mid v \in [\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*], \boldsymbol{\beta} \in [0, \infty)^{k-1} \right\}. \end{aligned} \tag{4.4}$$

For all $f \in \tilde{\mathcal{B}}_k$, we have that $f(\hat{\alpha}_{x|y}^*) = (1 - \hat{\alpha}_{x|y}^*)$ and $f(\hat{\alpha}_{y|x}^*) = \hat{\alpha}_{y|x}^*$, and each f is only defined on the interval $[\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*]$. Letting $\tilde{\lambda}(\cdot; \boldsymbol{\beta}) \in \tilde{\mathcal{B}}_k$ represent a form of the ADF for $w \in [\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*]$, the techniques introduced in Sects. 4.1 and 4.2 can be used to

obtain estimates of the coefficient vectors, which we denote $\hat{\boldsymbol{\beta}}_{CL2}$ and $\hat{\boldsymbol{\beta}}_{PR2}$, respectively. The resulting estimators for λ are

$$\hat{\lambda}_{CL2}(w) = \begin{cases} \tilde{\lambda}(w; \hat{\boldsymbol{\beta}}_{CL2}) & \text{for } w \in [\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*], \\ \max(w, 1 - w) & \text{for } w \in [0, \hat{\alpha}_{x|y}^*) \cup (\hat{\alpha}_{y|x}^*, 1], \end{cases}$$

with $\hat{\lambda}_{PR2}$ defined analogously. We lastly define the discontinuous estimator $\hat{\lambda}_{H2}$ as

$$\hat{\lambda}_{H2} := \begin{cases} \hat{\lambda}_H(w) & \text{for } w \in [\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*], \\ \max(w, 1 - w) & \text{for } w \in [0, \hat{\alpha}_{x|y}^*) \cup (\hat{\alpha}_{y|x}^*, 1]. \end{cases}$$

This is obtained by combining the pointwise Hill estimator with the information provided by the estimates $\hat{\alpha}_{x|y}^*$, $\hat{\alpha}_{y|x}^*$. Illustrations of all the estimators discussed in this section, as well as in Sect. 2, can be found in the Supplementary Material.

We note that estimation of the parameters, $\alpha_{x|y}$, $\alpha_{y|x}$, is subject to uncertainty and will not give a perfect representation of $\alpha_{x|y}^*$ and $\alpha_{y|x}^*$. However, we believe it is worth exploring the quality of the resulting estimates, and find in Sect. 5 that they generally provide improvement over estimators that do not exploit this link. We also note the potential for independent estimates of λ to be used for improving estimation of $\alpha_{x|y}$ and $\alpha_{y|x}$ in the conditional extremes model, but leave exploration of this to future work.

4.4 Incorporating theoretical properties

All estimators introduced so far are not required to satisfy the property of λ introduced in Proposition 3.1, or the lower bound on the ADF discussed in Sect. 1. Furthermore, the estimator $\hat{\lambda}_H$ is also not guaranteed to satisfy the endpoint conditions, i.e., $\lambda(0) = \lambda(1) = 1$.

There are several techniques one could use to impose these properties. For instance, one could incorporate penalty terms to objective functions $\mathcal{L}_C(\boldsymbol{\beta})$ or $S(\boldsymbol{\beta})$ to penalise for cases when the conditions of Proposition 3.1, or the ADF lower bound, are not satisfied. Alternatively, one could also impose the theoretical properties of the ADF in post-processing steps; such a procedure can be applied regardless of whether the original estimator is smooth or local. In unreported results, we considered penalising the objective function $\mathcal{L}_C(\boldsymbol{\beta})$ for violations of ADF properties, with larger penalties for greater violations. This approach allows for simpler optimisation of the objective function than imposing a very strong penalty for any violation, but introduces the choice of a penalty parameter and does not guarantee the properties are fully satisfied. We therefore opt instead to apply a post-processing procedure to each of the ADF estimators; this has the added advantage of also being applicable to the pointwise Hill estimators.

For any estimator $\hat{\lambda}_-$, assume that the set $\{w \in [0, 1] \mid \hat{\lambda}_-(w) < \max(w, 1 - w)\}$ is non-empty. To ensure the ADF is bounded from below, and satisfies the endpoint conditions $\lambda(0) = \lambda(1) = 1$, we set

$$\hat{\lambda}_-(w) = \begin{cases} \max \{ \hat{\lambda}_-(w), \max(w, 1 - w) \} & \text{for } w \in (0, 1), \\ 1 & \text{for } w \in \{0, 1\}. \end{cases}$$

Next, we ensure the conditions outlined in Proposition 3.1 are satisfied. Define the angular sets $\mathscr{W}^{\leq 0.5} = (w_1^{\leq 0.5}, w_2^{\leq 0.5}, \dots, w_{(m-1)/2}^{\leq 0.5}) := \{i/(m-1) : i = (m-1)/2, (m-1)/2 - 1, \dots, 0\} \subset \mathscr{W}$ and $\mathscr{W}^{\geq 0.5} = (w_1^{\geq 0.5}, w_2^{\geq 0.5}, \dots, w_{(m-1)/2}^{\geq 0.5}) := \{i/(m-1) : i = (m-1)/2, (m-1)/2 + 1, \dots, m-1\} \subset \mathscr{W}$. We propose the following algorithm.

Algorithm 1 Algorithm for imposing Property 3.1.

```

for  $i \leftarrow 2$  to  $(m-1)/2$  do
  if  $w_{i-1}^{\leq 0.5} / \hat{\lambda}_-(w_{i-1}^{\leq 0.5}) < w_i^{\leq 0.5} / \hat{\lambda}_-(w_i^{\leq 0.5})$  then
    | set  $\hat{\lambda}_-(w_i^{\leq 0.5}) := w_i^{\leq 0.5} \hat{\lambda}_-(w_{i-1}^{\leq 0.5}) / w_{i-1}^{\leq 0.5}$ ;
  end
  if  $(1 - w_{i-1}^{\geq 0.5}) / \hat{\lambda}_-(w_{i-1}^{\geq 0.5}) < (1 - w_i^{\geq 0.5}) / \hat{\lambda}_-(w_i^{\geq 0.5})$  then
    | set  $\hat{\lambda}_-(w_i^{\geq 0.5}) := (1 - w_i^{\geq 0.5}) \hat{\lambda}_-(w_{i-1}^{\geq 0.5}) / (1 - w_{i-1}^{\geq 0.5})$ ;
  end
end

```

This ensures the processed estimator $\hat{\lambda}_-$ satisfies the conditions of Properties 3.1 for $w \in \mathscr{W}$; this is the finite window that we use to represent $[0, 1]$ in practice. This processing is applied to all pointwise and novel estimators, i.e., $\hat{\lambda}_H, \hat{\lambda}_{CL}, \hat{\lambda}_{PR}, \hat{\lambda}_{H2}, \hat{\lambda}_{CL2}$ and $\hat{\lambda}_{PR2}$, ensuring that the ADF estimates from each approach are always theoretically valid. In practice, we found that imposing these theoretical results also improved estimation quality within the resulting ADF estimates, both in terms of bias and variance. An illustration of the processing procedure is given in the Supplementary Material.

4.5 Tuning parameter selection

Prior to using any of the ADF estimators introduced in this section, we are required to select at least one tuning parameter. For the probability values required by the estimators introduced in Sects. 4.1 and 4.2, we set $q = 0.90, \{q_j\}_{j=1}^h := \{0.87 + (j-1) \times 0.002\}_{j=1}^h$ and $p_j = q_j + 0.05$ for $j = 1, \dots, h$, with $h = 31$. These values were chosen to evaluate whether the resulting estimators improve upon the base estimator $\hat{\lambda}_H$ using (approximately) the same amount of tail information in all cases. We tested a range of probabilities for both estimators and found that the ADF estimates were not massively sensitive to these across different extremal dependence structures. For example, for $\hat{\lambda}_{CL}$, a lower q resulted in mild improvements for asymptotically independent copulas, while simultaneously worsening the quality of ADF estimates for asymptotically dependent examples, while a higher q led to higher variance.

For the angular interval \mathscr{W} , we set $m = 1,000$, i.e., $\mathscr{W} = \{0, 0.001, 0.002, \dots, 0.999, 1\}$. This set was sufficient to ensure a high degree of smoothness in the resulting ADF estimates without too high a computational burden.

For each of the novel estimators (except $\hat{\lambda}_{H2}$), we must also specify the degree $k \in \mathbb{N}$ for the polynomial families described by Eqs. 4.1 and 4.4. In the case of the Pickands' dependence function, studies have found that higher values of k are preferable for very strong positive dependence, while the opposite is true for weak dependence (Marcon et al. 2017; Vettori et al. 2018). We prefer to select a single value of k that performs well across a range of dependence structures, while minimising the computational burden; this avoids the need to select this parameter when obtaining ADF estimates in practice.

To achieve this objective, we estimated the root mean integrated squared error (RMISE), as defined in Sect. 5.1, of the estimators $\hat{\lambda}_{CL}$ and $\hat{\lambda}_{PR}$ with $k = 4, \dots, 11$ using 200 samples from two Gaussian copula examples, corresponding to strong ($\rho = 0.9$) and weak ($\rho = 0.1$) positive dependence. Assessment of how the RMISE estimates vary over k for both estimators suggests that $k = 7$ is sufficient to accurately capture both dependence structures. The full results can be found in the Supplementary Material. We remark that this approach for selecting k is somewhat ad hoc, and in practice, one could employ various diagnostic tools, such as the tool discussed in Sect. 6.3, to select k on a case-by-case basis.

For each of the 'combined' estimators in Sect. 4.3, we take the same tuning parameters as for the 'non-combined' counterpart, since the combined estimators have near identical formulations only defined on a subset of $[0, 1]$. For example, the empirical 90% threshold of the min-projection is used for both $\hat{\lambda}_H$ and $\hat{\lambda}_{H2}$. Finally, when estimating the conditional extremes parameters, empirical 90% conditioning thresholds are used.

5 Simulation study

5.1 Overview

In this section, we use simulation to compare the estimators proposed in Sect. 4 to the existing techniques described in Sect. 2. For the comparison, we introduce nine copula examples, representing a wide variety of extremal dependence structures, and encompassing both extremal dependence regimes.

The first three examples are from the bivariate Gaussian distribution, for which the dependence is characterised by the correlation coefficient $\rho \in [-1, 1]$. We consider $\rho \in \{-0.6, 0.1, 0.6\}$, resulting in data structures exhibiting medium negative, weak positive, and medium positive dependence, respectively. Note that in the case of $\rho = -0.6$, the choice of exponential margins will hide the dependence structure (Keef et al. 2013a; Nolde and Wadsworth 2022).

For the next two examples, we consider the bivariate extreme value copula with logistic and asymmetric logistic families (Gumbel 1960; Tawn 1988). In both cases, the dependence is characterised by the parameter $r \in (0, 1]$; we set $r = 0.8$, corresponding to weak positive extremal dependence. For the asymmetric logistic family, we also require two asymmetry parameters $(k_1, k_2) \in [0, 1]^2$, which we fix to be $(k_1, k_2) = (0.3, 0.7)$, resulting in a mixture structure.

We next consider the inverted bivariate extreme value copula (Ledford and Tawn 1997) for the logistic and asymmetric logistic families, with the dependence again characterised by the parameters r and (r, k_1, k_2) , respectively. We set $r = 0.4$, corresponding to moderate positive dependence, and again fix $(k_1, k_2) = (0.3, 0.7)$. Note that for this copula, the model described in Eq. 1.5 is exact: see Wadsworth and Tawn (2013).

Lastly, we consider the bivariate student t copula, for which dependence is characterised by the correlation coefficient $\rho \in [-1, 1]$ and the degrees of freedom $\nu > 0$. We consider $\rho = 0.8, \nu = 2$ and $\rho = 0.2, \nu = 5$, corresponding to strong and weak positive dependence.

Illustrations of the true ADFs for each copula are given in Fig. 4, showing a range of extremal dependence structures. For examples where the ADF equals the lower bound, the copula exhibits asymptotic dependence. While the fifth copula exhibits

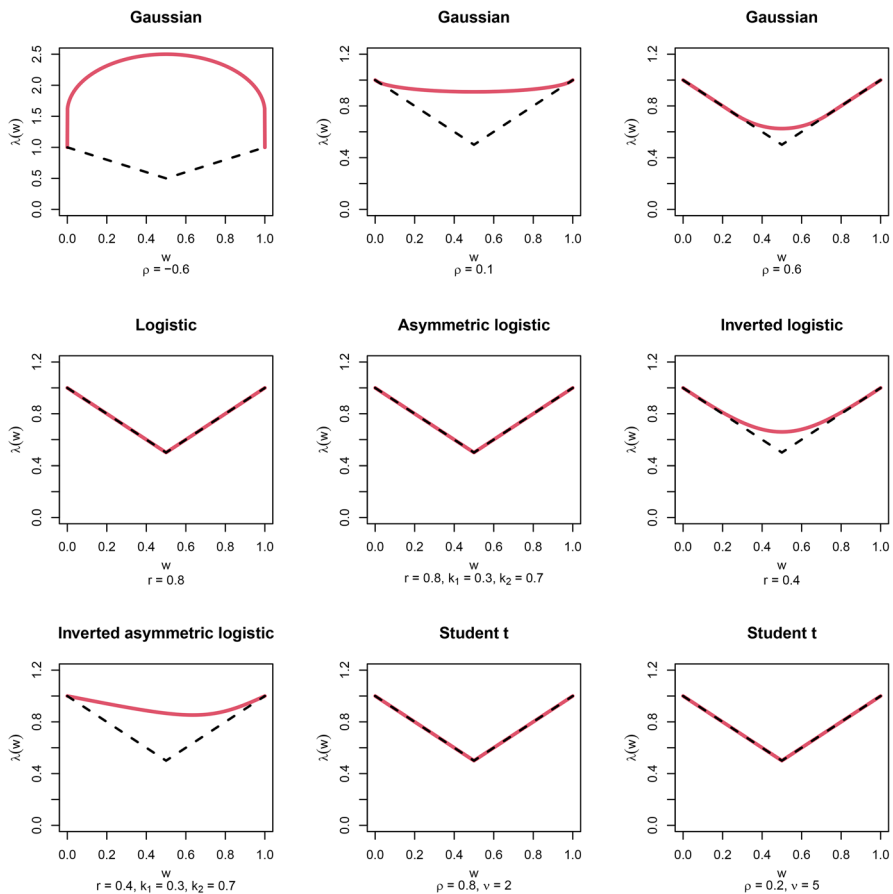


Fig. 4 True ADFs (in red) for each copula introduced in Sect. 5.1, along with the theoretical lower bound (black dotted line)

asymmetric dependence, the limiting ADF is symmetric; the same is not true for its inverted counterpart.

To evaluate estimator performance, we use the RMISE

$$\text{RMISE}(\hat{\lambda}_-) = \left(\mathbb{E} \left[\int_0^1 \{ \hat{\lambda}_-(w) - \lambda(w) \}^2 dw \right] \right)^{1/2},$$

where $\hat{\lambda}_-$ denotes an arbitrary estimator. Simple rearrangement shows that this metric is equal to the square root of the sum of integrated squared bias (ISB) and integrated variance (IV) (Gentle, 2002), i.e.,

$$\text{RMISE}(\hat{\lambda}_-) = \left(\underbrace{\int_0^1 [\mathbb{E}\{\hat{\lambda}_-(w)\} - \lambda(w)]^2 dw}_{\text{ISB}} + \underbrace{\int_0^1 \mathbb{E} \left([\mathbb{E}\{\hat{\lambda}_-(w)\} - \hat{\lambda}_-(w)]^2 \right) dw}_{\text{IV}} \right)^{1/2}.$$

Therefore, the RMISE summarises the quality of an estimator in terms of both bias and variance, and can be used as a means to compare different estimators.

5.2 Results

For the copulas described in Sect. 5.1, data from each copula example was simulated on standard exponential margins with a sample size of $n = 10,000$, and the integrated squared error (ISE) of each estimator was approximated for 1,000 samples using the trapezoidal rule; see the Supplementary Material for further information. The square root of the mean of these estimates was then computed, resulting in a Monte–Carlo estimate of the RMISE.

The RMISE estimates for each estimator and copula combination are shown in Table 1. Tables containing the corresponding Monte–Carlo error, ISB and IV values can be found in the Supplementary Material, along with root mean squared error estimates for individual rays $w \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. For each estimator, the bias varies significantly across the different copulas. However, in the majority of cases, the bias/variance are similar across most of the estimators. We remark that the magnitudes of the Monte–Carlo errors, as reported in the Supplementary Material, are small enough such that the ordering of RMISE estimates in Table 1 is likely to be a true reflection of the relative performance for each estimator.

While no estimator consistently outperforms the others, $\hat{\lambda}_{CL2}$ and $\hat{\lambda}_{ST}$ tend to have lower RMISE, ISB and IV values, on average. This is especially the case when comparing to the base estimator $\hat{\lambda}_H$. Furthermore, the ‘combined’ estimators outperform their non-combined counterparts in many cases, suggesting that incorporating parameter estimates from the conditional extremes model can reduce bias and variance. The Gaussian copula with $\rho = -0.6$ has much higher RMISE values, indicating that none of the estimators capture negative dependence well, though this is in part due to the choice of exponential margins.

Table 1 RMISE values (multiplied by 100) for each estimator and copula combination

Copula	$\hat{\lambda}_H$	$\hat{\lambda}_{CL}$	$\hat{\lambda}_{PR}$	$\hat{\lambda}_{H2}$	$\hat{\lambda}_{CL2}$	$\hat{\lambda}_{PR2}$	$\hat{\lambda}_{ST}$
Copula 1	61.2	61.3	66.2	61.4	61.9	66.7	63.7
Copula 2	3.44	3.36	3.67	3.41	3.35	3.66	2.95
Copula 3	3.43	3.46	3.83	3.21	3.22	3.57	1.09
Copula 4	4.58	4.71	6.89	4.24	4.24	6.17	2.77
Copula 5	14	14.1	17.1	14	14.1	17.1	12.1
Copula 6	2.05	2	2.18	1.78	1.75	1.92	2.12
Copula 7	2.79	2.68	2.93	2.77	2.69	2.93	3.96
Copula 8	1.04	1.05	1.44	0.562	0.535	0.72	1.87
Copula 9	11.9	12	14.9	11.9	12	14.9	11.1

Smallest RMISE values in each row are highlighted in bold, with values reported to 3 significant figures. Copulas 1-3 refer to the Gaussian copula with $\rho = -0.6$, $\rho = 0.1$ and $\rho = 0.6$, respectively. Copulas 4-7 refer to the logistic, asymmetric logistic, inverted logistic, inverted asymmetric logistic copulas, respectively. Copulas 8-9 refer to the student t copula with parameters $\rho = 0.8$, $\nu = 2$ and $\rho = 0.2$, $\nu = 5$, respectively

We note that while the true ADFs are identical for the asymptotically dependent logistic, asymmetric logistic, and student t copulas, the corresponding RMISE values in Table 1 vary significantly. Notably, the asymmetric logistic, which possesses a complex asymmetric structure, has substantially higher RMISE values compared to the other asymptotically dependent copulas. We suspect these disparities arise at finite levels due to the different rates of convergence to the limiting ADF in Eq. 1.5, alongside the fact many multivariate extreme value models perform poorly in the case of asymmetric dependence (Tendijck et al. 2021).

Overall, these results indicate that no one estimator is preferable across all extremal dependence structures. However, we suggest using the estimators $\hat{\lambda}_{CL2}$ and $\hat{\lambda}_{ST}$ since, on average, these appeared to result in less bias and variance. The form of extremal dependence appears to affect the performance of both of these estimators; since this is often difficult to quantify a priori, we suggest using both estimators and evaluating relative performance via diagnostics, as we do in Sect. 6.

6 Case study

6.1 Overview

Understanding the probability of observing extreme river flow events (i.e., floods) at multiple sites simultaneously is important in a variety of sectors, including insurance (Keef et al. 2013b; Quinn et al. 2019; Rohrbeck and Cooley 2021) and environmental management (Lamb et al. 2010; Gouldby et al. 2017). Valid risk assessments therefore require accurate evaluation of the extremal dependence between different sites.

In this section, we estimate the ADF of river flow data sets obtained from gauges in the north of England, UK, which can be subsequently used to construct bivariate

return curves. Daily average flow values (m^3/s) at six river gauge locations on different rivers were considered. The gauge sites are illustrated in Fig. 5. For each location, data is available between May 1993 and September 2020; however, we only consider dates where a measurement is available for every location. To avoid seasonality, we consider the interval October-March only; from our analysis, it appears that the highest daily flow values are observed in this period. This results in $n = 4,659$ data points for each site. Plots of the daily flow time series can be found in the Supplementary Material; these plots suggest that an assumption of stationarity is reasonable for the extremes of each data set.

We fix the site on the river Lune to be our reference site and consider the extremal dependence between this and all other gauges. We first estimate the extremal dependence measure χ and the coefficient η using the upper 10% of the corresponding joint tails. Both χ and η are limiting values; however, in practice, we are

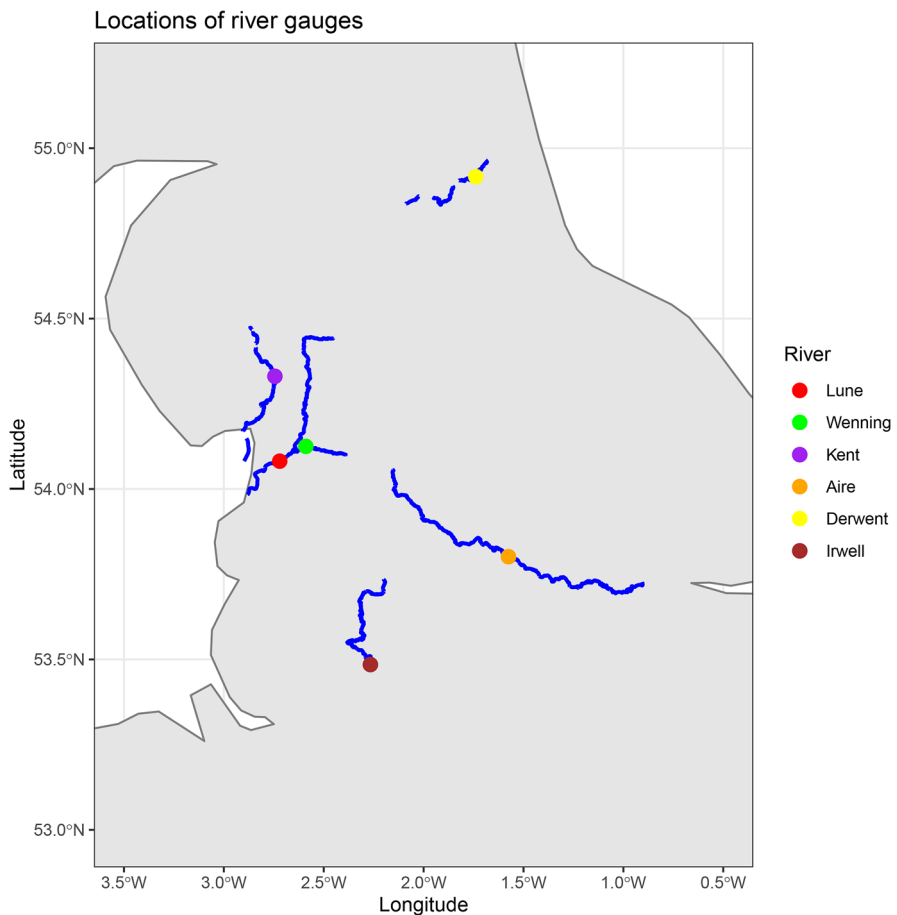


Fig. 5 Locations of river gauges in the north of England, UK. Individual rivers illustrated in blue alongside the corresponding gauge locations

unable to evaluate such limits without a closed form for the joint distribution. We therefore calculate these values empirically. Taking χ , for example, an estimate is $\hat{\chi}_q = \hat{\Pr}(X > \hat{x}_q, Y > \hat{y}_q) / \hat{\Pr}(X > \hat{x}_q)$, where $\hat{\Pr}(\cdot)$ denotes an empirical probability estimate and \hat{x}_q and \hat{y}_q denote empirical q quantile estimates for the variables X and Y , respectively, and q is some value close to 1. Specifically, we take $q = 0.9$. In practice, we are unlikely to observe $\chi = \hat{\chi}_q$, even at the most extreme quantile levels, i.e., as $q \rightarrow 1$. This can be problematic when trying to quantify the form of extremal dependence, since $\hat{\chi}_q > 0$ may arise for asymptotically independent data sets (for example). Therefore, the estimated coefficients should act only as a rough guide for this quantification.

The dependence measure estimates and 95% confidence intervals are shown in Fig. 6 as a function of distance from the reference site. Here and throughout, all confidence intervals are obtained via block bootstrapping with block size $b = 40$; this value appears appropriate to account for the varying degrees of temporal dependence observed across the six gauge sites. These estimates suggest that asymptotic independence may exist for at least three of the site pairings; therefore, modelling techniques based on bivariate regular variation would likely fail to capture the observed extremal dependencies in this scenario.

6.2 ADF estimation

We transform each marginal data set to exponential margins using the semi-parametric approach of Coles and Tawn (1991), whereby a generalised Pareto distribution is fitted to the upper tail and the body is modelled empirically. The generalised Pareto distribution thresholds are selected using the technique proposed in Murphy et al. (2024). In spite of the data violating the independence assumption, diagnostic plots found in the Supplementary Material indicate reasonable model fits. Since our results from Sect. 5 suggest that the estimators $\hat{\lambda}_{CL2}$ and $\hat{\lambda}_{ST}$ perform best overall, we used these, alongside the base estimator $\hat{\lambda}_H$, to estimate the ADF for each combination of the reference gauge and the other five gauges. The resulting ADF estimates can be found in Fig. 7.

One can observe contrasting shapes across the different pairs of gauges, illustrating the variety in the observed extremal dependence structures. These results illustrate that on the whole, the estimator $\hat{\lambda}_{CL2}$ is very much a smoothed version of $\hat{\lambda}_H$ on the interval $(\hat{\alpha}_{x|y}^*, \hat{\alpha}_{y|x}^*)$, owing to the form of likelihood function used.

6.3 Assessing goodness of fit for ADF estimates

Recall that, from Eq. 2.1, we have $T_w^* \sim \text{Exp}(\lambda(w))$ as $u_w \rightarrow \infty$ for all $w \in [0, 1]$. We exploit this result to assess the goodness of fit for ADF estimates.

Let $\hat{\lambda}(w), w \in [0, 1]$, denote an estimated ADF obtained using the sample. Given $w \in [0, 1]$, let u_w denote some high empirical quantile from the sample \mathbf{t}_w , and consider the sample \mathbf{t}_w^* , with \mathbf{t}_w and \mathbf{t}_w^* defined as in Sect. 4.1. If \mathbf{t}_w^* is indeed a sample from an $\text{Exp}(\hat{\lambda}(w))$ distribution, we would expect good agreement between

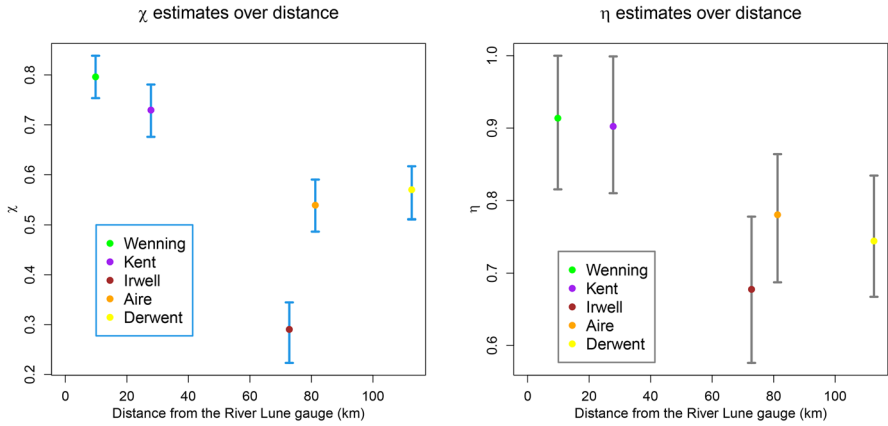


Fig. 6 Estimated dependence coefficients as a function of distance from the Lune gauge, with 95% point-wise confidence intervals given by the shaded regions. Left: Estimates of χ (blue). Right: Estimates of η (grey)

empirical and model quantiles. Analogously, $\hat{\lambda}(w)\mathbf{t}_w^*$ should represent a sample from an $\text{Exp}(1)$ distribution if $T_w^* \sim \text{Exp}(\lambda(w))$, independent of the ray w .

We use these results to derive local and global diagnostics for the ADF. First, for a subset of rays $w \in [0, 1]$, corresponding to a range of joint survival regions,

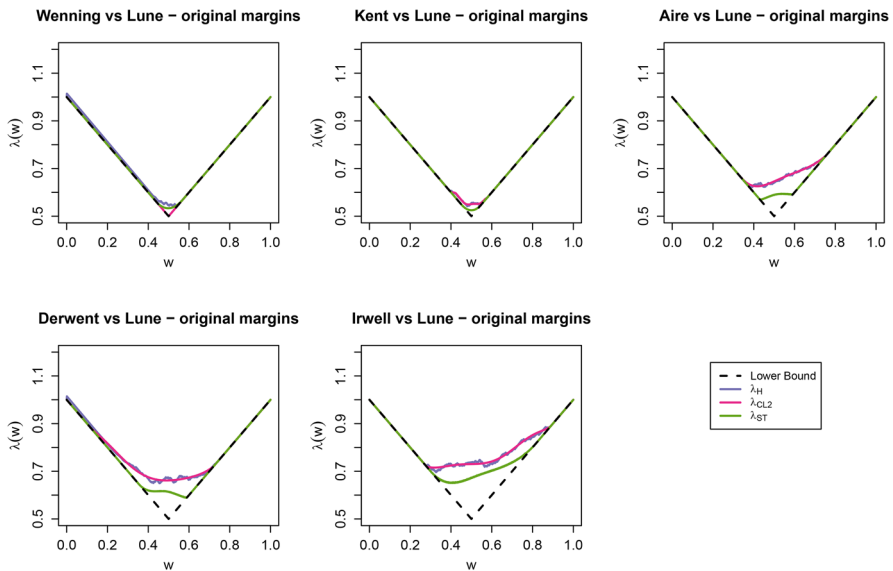


Fig. 7 ADF estimates for each pair of gauge sites. The purple, pink and green lines represent the estimates from $\hat{\lambda}_H$, $\hat{\lambda}_{CL2}$ and $\hat{\lambda}_{ST}$, respectively, with the theoretical lower bound denoted by the black dotted lines

let $n_w = |\mathbf{t}_w^*|$ and $\mathbf{t}_{w(j)}^*$ denote the j -th order statistic of \mathbf{t}_w^* and consider the set of pairs

$$\left\{ \left(-\log(1 - j/(n_w + 1)) / \hat{\lambda}(w), \mathbf{t}_{w(j)}^* \right) : j = 1, \dots, n_w \right\}.$$

With u_w fixed to be the 90% empirical quantile of \mathbf{t}_w , quantile-quantile (QQ) plots for five individual rays, $w \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, are given in the first five panels Fig. 8 for the third pair of gauges and the $\hat{\lambda}_{CL2}$ estimator. Uncertainty intervals are obtained via block bootstrapping on the set \mathbf{t}_w^* , i.e., the order statistics. We acknowledge a deficiency of this scheme that all sampled quantiles will be bounded by the interval $[\mathbf{t}_{w(1)}^*, \mathbf{t}_{w(n_w)}^*]$. However, alternative uncertainty quantification approaches, such as parametric bootstraps or the use of beta quantiles for uniform order statistics, would fail to account for the observed temporal dependence within the data.

For the global diagnostic, we propose the following: for each $i \in \{1, \dots, n\}$, define the corresponding angular observation $w_i := x_i / (x_i + y_i)$ and let u_{w_i} denote the 90% empirical quantile of \mathbf{t}_{w_i} . Randomly sample one point t_i^* from the set $\mathbf{t}_{w_i}^*$ and set $e_i := \lambda(w_i)(t_i^* - u_{w_i})$; repeating this process over all $i \in \{1, \dots, n\}$, we obtain the set $\mathcal{E} := \{e_i \mid i \in \{1, \dots, n\}\}$. We then consider the set of pairs

$$\left\{ \left(-\log(1 - j/(n + 1)), e_{(j)} \right) : j = 1, \dots, n \right\},$$

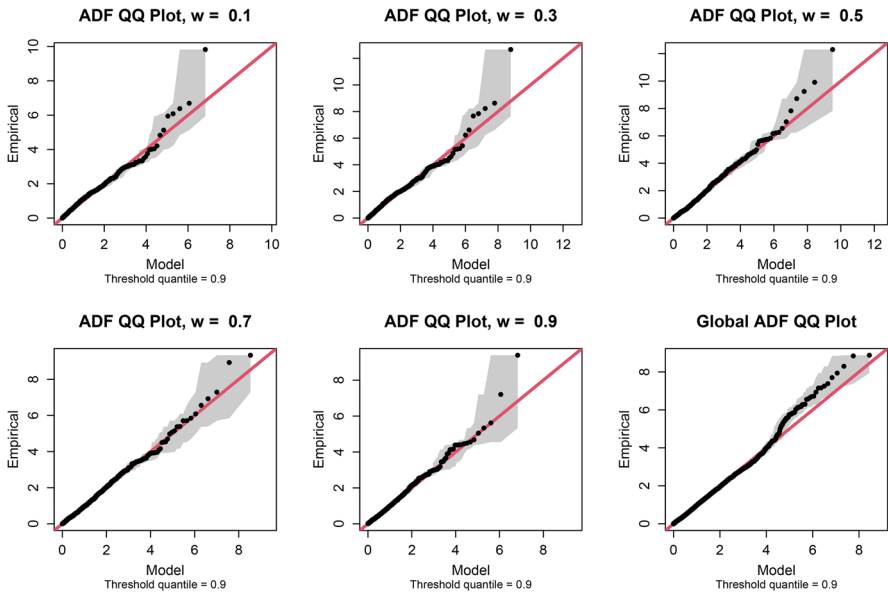


Fig. 8 Local and global ADF QQ plots for the third pair of gauges, using the ADF estimate obtained via $\hat{\lambda}_{CL2}$. Estimates given in black, with 95% pointwise tolerance intervals represented by the grey shaded regions. Red lines correspond to the $y = x$ line

with $e_{(j)}$ denoting the j -th order statistic of \mathcal{E} ; this comparison provides an overall summary for the quality of model fit across all angles. The corresponding QQ plot for the third pair of gauges and the $\hat{\lambda}_{CL2}$ estimator is given in the bottom right panel of Fig. 8, with uncertainty bounds obtained via block bootstrapping on set \mathcal{E} . This diagnostic possesses a degree of stochasticity due to sampling from the set $\mathbf{t}_{w_i}^*$. We can check the impact of this on our impression of the fit by considering a few different random seeds; see the Supplementary Material for further details.

On the whole, the estimated exponential quantiles appear in good agreement with the observed quantiles, indicating the underlying ADF estimate accurately captures the tail behaviour for each min-projection variable. Analogous plots for $\hat{\lambda}_H$ and $\hat{\lambda}_{ST}$ are given in the Supplementary Material. Similar plots were obtained when the other pairs of gauges were considered.

6.4 Estimating return curves

To quantify the risk of joint flooding events across sites, we follow Murphy-Barltrop et al. (2023) and use the ADF to estimate a bivariate risk measure known as a return curve, $RC(p)$, as defined in Sect. 1. This measure is a direct bivariate extension of a return level, which is commonly used to quantify risk in the univariate setting (Coles 2001). Taking probability values p close to zero gives a summary of the joint extremal dependence, thus allowing for comparison across different data structures. In the context of extreme river flows, return curves can be used to evaluate at which sites joint extremes (floods) are more/less likely to occur. For illustration, we fix p to correspond to a 5 year return period, i.e., $p = 1/(5n_y)$, where n_y is the average number of points observed in a given year (Brunner et al. 2016). Excluding missing observations, we have 28 years of data, hence the resulting curve should be within the range of data whilst simultaneously representing the joint tail. The resulting return curve estimates for each ADF estimator and pair of gauge sites can be found in Fig. 9.

There is generally good agreement among the estimated curves. The almost-square shapes of the estimates for the first two pairs of gauges indicate higher likelihoods of observing simultaneous flood events at the corresponding gauge sites; this is as expected given the close spatial proximity of these sites. In all cases, the curves derived via $\hat{\lambda}_H$ are quite rough and unrealistic, and are subsequently ignored. To assess the goodness of fit of the remaining return curve estimates, we consider the first and fifth examples and apply the diagnostic introduced in Murphy-Barltrop et al. (2023). Our results suggest good quality model fits for both of the estimates obtained using $\hat{\lambda}_{CL2}$ and $\hat{\lambda}_{ST}$, though with potentially a slight preference for the estimate based on $\hat{\lambda}_{CL2}$ for the fifth pair. Furthermore, we also obtain 95% return curve confidence intervals for these examples. The resulting plots illustrating the diagnostics and confidence intervals, along with a brief explanation of the diagnostic tool, are given in the Supplementary Material.

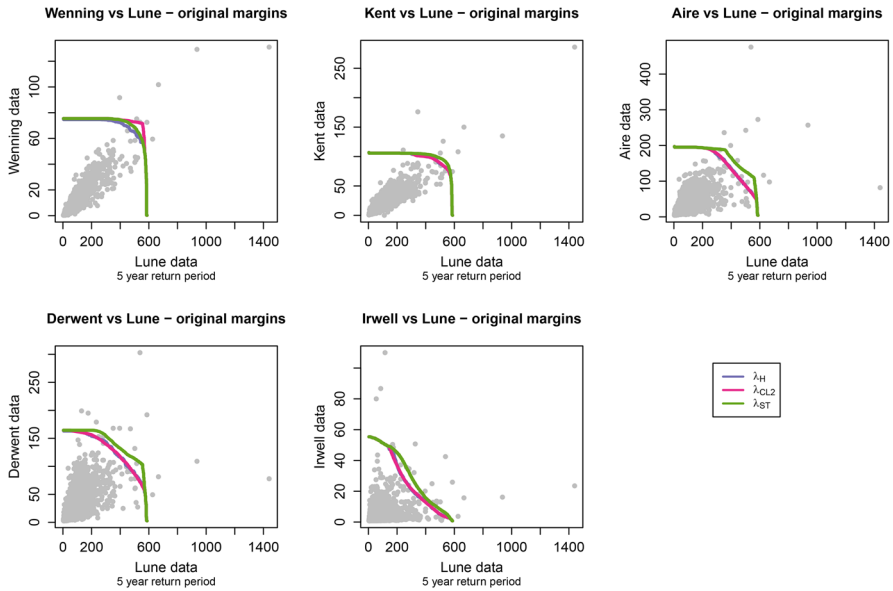


Fig. 9 Estimated 5-year return curves (on original margins) for each pair of gauges. The purple, pink and green lines represent the curve estimates from $\hat{\lambda}_H$, $\hat{\lambda}_{CL2}$ and $\hat{\lambda}_{ST}$, respectively

7 Discussion

We have introduced a range of novel global estimators for the ADF, as detailed in Sect. 4. We compared these estimators to existing techniques through a systematic simulation study and found our novel estimators to be competitive in many cases. In particular, the estimators derived via the composite likelihood approach of Sect. 4.1, alongside the estimator of Simpson and Tawn (2022), appear to have lower bias and variance, on average, compared to alternative estimation techniques. We also applied ADF estimation techniques to a range of river flow data sets, and obtained estimates of return curves for each data set. The results suggest that our estimation procedures are able to accurately capture the range of extremal dependence structures exhibited in the data.

From Sect. 5, one can observe that the ‘combined’ estimators proposed in Sect. 4.3 outperform their ‘uncombined’ counterparts in the majority of instances. This indicates that incorporating the knowledge obtained from the conditional extremes parameters leads to improvements in ADF estimates. Furthermore, in most cases, ADF estimates obtained via approximations of the set ∂G appeared to have lower bias compared to alternative estimation techniques. More generally, these results suggest that inferential techniques that exploit the results of Nolde and Wadsworth (2022) are superior to techniques which do not. Estimation of ∂G , and its impact on estimation of other extremal dependence properties, represents an important line of research.

As noted in Sect. 1, few applications of the modelling framework described in Eq. 1.5 exist, even though this model offers advantages over the widely used approach of Heffernan and Tawn (2004) when evaluating joint tail probabilities. Inference via the ADF ensures consistency in extremal dependence properties, and one can obtain accurate estimates of certain risk measures, such as return curves.

For each of the existing and novel estimators introduced in Sects. 2 and 4, we were required to select quantile levels, which is equivalent to selecting thresholds of the min-projection. With the exception of $\hat{\lambda}_{ST}$, similar quantile levels were considered for each estimator so as to provide some degree of comparability. However, due to variation in estimation procedures, we acknowledge that the selected quantile levels are not readily comparable since the quantity of joint tail data used for estimation varies across different estimators. Moreover, as noted in Sect. 4.5, trying to select ‘optimal’ quantile levels appears a fruitless exercise since the performance of each estimator does not appear to alter much across different quantile levels.

As noted in Sect. 4.5, our proposed estimators require selection of several tuning parameters. For all practical applications, we recommend this selection is done using a combination of the diagnostic tools outlined in Sect. 6, since it is unlikely that one set of tuning parameters will be appropriate across all observed dependence structures and sample sizes.

Finally, we acknowledge the lack of theoretical treatment for the proposed ADF estimators which is an important consideration for understanding properties of the methodology. However, theoretical results of this form typically require in-depth analyses and strict assumptions, which themselves may be hard to verify, whilst in practice one can only ever look at diagnostics obtained from the data. We have therefore opted for a more practical treatment of the proposed estimators.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10687-024-00490-4>.

Acknowledgements This paper is based on work partly completed while Callum Murphy-Barltrop was part of the EPSRC funded STOR-i centre for doctoral training (EP/L015692/1). We are grateful to the referee and editor for constructive comments that have improved this article.

Author Contributions All authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by Callum Murphy-Barltrop. The first draft of the manuscript was written by Callum Murphy-Barltrop and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL. This work was supported by EPSRC grant numbers EP/L015692/1 and EP/X010449/1.

Data Availability The data sets analysed in Sect. 6 are freely available online from the National River Flow Archive (National River Flow Archive 2022). The gauge ID numbers for the six considered stations are as follows: River Lune - 72004, River Wenning - 72009, River Kent - 73012, River Irwell - 69002, River Aire - 27028, River Derwent - 23007. Date downloaded: October 24th, 2022.

Declarations

Ethical Approval Not Applicable

Conflict of interest The authors declare no Conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Brunner, M.I., Seibert, J., Favre, A.C.: Bivariate return periods and their importance for flood peak and volume estimation. *Wiley Interdiscip. Rev. Water* **3**, 819–833 (2016)
- Castro-Camilo, D., de Carvalho, M., Wadsworth, J.: Time-varying extreme value dependence with application to leading European stock markets. *Ann. Appl. Stat.* **12**, 283–309 (2018)
- Coles, S.: An introduction to statistical modeling of extreme values. Springer, London (2001)
- Coles, S., Pauli, F.: Models and inference for uncertainty in extremal dependence. *Biometrika* **89**, 183–196 (2002)
- Coles, S.G., Tawn, J.A.: Modelling extreme multivariate events. *J. R. Stat. Soc. Ser. B Stat Methodol.* **53**, 377–392 (1991)
- Cormier, E., Genest, C., Nešlehová, J.G.: Using B-splines for nonparametric inference on bivariate extreme-value copulas. *Extremes* **17**, 633–659 (2014)
- Davison, A.C., Smith, R.L.: Models for exceedances over high thresholds. *J. R. Stat. Soc. Ser. B Stat Methodol.* **52**, 393–425 (1990)
- de Carvalho, M., Davison, A.C.: Spectral density ratio models for multivariate extremes. *J. Am. Stat. Assoc.* **109**, 764–776 (2014)
- Eastoe, E.F., Heffernan, J.E., Tawn, J.A.: Nonparametric estimation of the spectral measure, and associated dependence measures, for multivariate extreme values using a limiting conditional representation. *Extremes* **17**, 25–43 (2014)
- Einmahl, J.H., Segers, J.: Maximum empirical likelihood estimation of the spectral measure of an extreme-value distribution. *Ann. Stat.* **37**, 2953–2989 (2009)
- Gentle, J.E.: *Elements of Computational Statistics*. Springer-Verlag (2002)
- Gouldby, B., Wyncoll, D., Panzeri, M., Franklin, M., Hunt, T., Hames, D., Tozer, N., Hawkes, P., Dornbusch, U., Pullen, T.: Multivariate extreme value modelling of sea conditions around the coast of England. *Proc. Inst. Civ. Eng. Marit. Eng.* **170**, 3–20 (2017)
- Guillotte, S., Perron, F.: Polynomial Pickands functions. *Bernoulli* **22**, 213–241 (2016)
- Gumbel, E.J.: Bivariate exponential distributions. *J. Am. Stat. Assoc.* **55**, 698–707 (1960)
- Hall, P., Tajvidi, N.: Distribution and dependence-function estimation for bivariate extreme-value distributions. *Bernoulli* **6**, 835–844 (2000)
- Haselsteiner, A.F., Coe, R.G., Manuel, L., Chai, W., Leira, B., Clarindo, G., Soares, C.G., Ásta Hannesdóttir, Dimitrov, N., Sander, A., Ohlendorf, J.H., Thoben, K.D., de Hauteclocque, G., Mackay, E., Jonathan, P., Qiao, C., Myers, A., Rode, A., Hildebrandt, A., Schmidt, B., Vanem, E., Huseby, A.B.: A benchmarking exercise for environmental contours. *Ocean Eng.* **236**, 1–29 (2021)
- Heffernan, J.E., Tawn, J.A.: A conditional approach for multivariate extreme values. *J. R. Stat. Soc. Ser. B Stat Methodol.* **66**, 497–546 (2004)
- Hill, B.M.: A simple general approach to inference about the tail of a distribution. *Ann. Stat.* **3**, 1163–1174 (1975)
- Huser, R., Wadsworth, J.L.: Modeling spatial processes with unknown extremal dependence class. *J. Am. Stat. Assoc.* **114**, 434–444 (2019)
- Joe, H.: *Multivariate Models and Multivariate Dependence Concepts*. Chapman and Hall/CRC (1997)

- Jonathan, P., Ewans, K., Flynn, J.: On the estimation of ocean engineering design contours. *J. Offshore Mech. Arct. Eng.* **136**, 1–8 (2014)
- Keef, C., Papastathopoulos, I., Tawn, J.A.: Estimation of the conditional distribution of a multivariate variable given that one of its components is large: Additional constraints for the Heffernan and Tawn model. *J. Multivar. Anal.* **115**, 396–404 (2013a)
- Keef, C., Tawn, J.A., Lamb, R.: Estimating the probability of widespread flood events. *Environmetrics* **24**, 13–21 (2013b)
- Lamb, R., Keef, C., Tawn, J., Laeger, S., Meadowcroft, I., Surendran, S., Dunning, P., Batstone, C.: A new method to assess the risk of local and widespread flooding on rivers and coasts. *Journal of Flood Risk Management* **3**, 323–336 (2010)
- Ledford, A.W., Tawn, J.A.: Statistics for near independence in multivariate extreme values. *Biometrika* **83**, 169–187 (1996)
- Ledford, A.W., Tawn, J.A.: Modelling dependence within joint tail regions. *J. R. Stat. Soc. Ser. B Stat Methodol.* **59**, 475–499 (1997)
- Liu, Y., Tawn, J.A.: Self-consistent estimation of conditional multivariate extreme value distributions. *J. Multivar. Anal.* **127**, 19–35 (2014)
- Marcon, G., Padoan, S.A., Antoniano-Villalobos, I.: Bayesian inference for the extremal dependence. *Electron. J. Stat.* **10**, 3310–3337 (2016)
- Marcon, G., Padoan, S.A., Naveau, P., Muliere, P., Segers, J.: Multivariate nonparametric estimation of the Pickands dependence function using Bernstein polynomials. *J. Stat. Plan. Inference* **183**, 1–17 (2017)
- Mhalla, L., de Carvalho, M., Chavez-Demoulin, V.: Regression-type models for extremal dependence. *Scand. J. Stat.* **46**, 1141–1167 (2019a)
- Mhalla, L., Opitz, T., Chavez-Demoulin, V.: Exceedance-based nonlinear regression of tail dependence. *Extremes* **22**, 523–552 (2019b)
- Murphy, C., Tawn, J.A., Varty, Z.: Automated threshold selection and associated inference uncertainty for univariate extremes. [arXiv:2310.17999](https://arxiv.org/abs/2310.17999) (2024)
- Murphy-Barttrop, C.J.R., Wadsworth, J.L.: Modelling non-stationarity in asymptotically independent extremes. *Comput. Stat. Data Anal.* **199**, 1–18 (2024)
- Murphy-Barttrop, C.J.R., Wadsworth, J.L., Eastoe, E.F.: New estimation methods for extremal bivariate return curves. *Environmetrics* **e2797**, 1–22 (2023)
- National River Flow Archive (2022). Daily mean river flow datasets. <https://nrfa.ceh.ac.uk/>. Accessed 11 Nov 2022
- Nolde, N.: Geometric interpretation of the residual dependence coefficient. *J. Multivar. Anal.* **123**, 85–95 (2014)
- Nolde, N., Wadsworth, J.L.: Linking representations for multivariate extremes via a limit set. *Adv. Appl. Probab.* **54**, 688–717 (2022)
- Pickands, J.: Multivariate extreme value distribution. *Proceedings 43th, Session of International Statistical Institution*, 1981 (1981)
- Quinn, N., Bates, P.D., Neal, J., Smith, A., Wing, O., Sampson, C., Smith, J., Heffernan, J.: The spatial dependence of flood hazard and risk in the United States. *Water Resour. Res.* **55**, 1890–1911 (2019)
- Ramos, A., Ledford, A.: A new class of models for bivariate joint tails. *J. R. Stat. Soc. Ser. B Stat Methodol.* **71**, 219–241 (2009)
- Resnick, S.I.: *Extreme Values. Regular Variation and Point Processes*. Springer, New York (1987)
- Resnick, S.: Hidden regular variation, second order regular variation and asymptotic independence. *Extremes* **5**, 303–336 (2002)
- Rohrbeck, C., Cooley, D.: Simulating flood event sets using extremal principal components. [arXiv:2106.00630](https://arxiv.org/abs/2106.00630) (2021)
- Ross, E., Astrup, O.C., Bitner-Gregersen, E., Bunn, N., Feld, G., Gouldby, B., Huseby, A., Liu, Y., Randell, D., Vanem, E., Jonathan, P.: On environmental contours for marine and coastal design. *Ocean Eng.* **195**, 106194 (2020)
- Simpson, E.S., Tawn, J.A.: Estimating the limiting shape of bivariate scaled sample clouds: with additional benefits of self-consistent inference for existing extremal dependence properties. [arXiv:2207.02626](https://arxiv.org/abs/2207.02626) (2022)
- Simpson, E.S., Wadsworth, J.L., Tawn, J.A.: Determining the dependence structure of multivariate extremes. *Biometrika* **107**, 513–532 (2020)
- Tawn, J.A.: Bivariate extreme value theory: models and estimation. *Biometrika* **75**, 397–415 (1988)
- Tendijck, S., Eastoe, E., Tawn, J., Randell, D., Jonathan, P.: Modeling the extremes of bivariate mixture distributions with application to oceanographic data. *J. Am. Stat. Assoc.* **0**, 1–12 (2021)
- Varin, C., Reid, N., Firth, D.: An overview of composite likelihood methods. *Stat. Sin.* **21**, 5–42 (2011)
- Vettori, S., Huser, R., Genton, M.G.: A comparison of dependence function estimators in multivariate extremes. *Stat. Comput.* **28**, 525–538 (2018)

- Wadsworth, J.L., Campbell, R.: Statistical inference for multivariate extremes via a geometric approach. *J. R. Stat. Soc. Series B: Statistical Methodology* qkae030 (2024)
- Wadsworth, J.L., Tawn, J.A.: A new representation for multivariate tail probabilities. *Bernoulli* **19**, 2689–2714 (2013)
- Wadsworth, J.L., Tawn, J.A., Davison, A.C., Elton, D.M.: Modelling across extremal dependence classes. *J. R. Stat. Soc. Ser. B Stat Methodol.* **79**, 149–175 (2017)
- Youngman, B.D.: Generalized additive models for exceedances of high thresholds with an application to return level estimation for U.S. wind gusts. *J. Am. Stat. Assoc.* **114**, 1865–1879 (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.