

Modelling non-stationary extremal dependence through a geometric approach

C. J. R. Murphy-Barltrop^{1,2*}, J. L. Wadsworth³, M. de Carvalho^{4,5}, and B. D. Youngman⁶

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

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

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

⁴School of Mathematics, University of Edinburgh EH9 3FD, UK

⁵CIDMA—Department of Mathematics, University of Aveiro, Portugal

⁶Department of Mathematics and Statistics, University of Exeter EX4 4QE, UK

*Correspondence to: callum.murphy-barltrop@tu-dresden.de

June 11, 2026

Abstract

Non-stationary extremal dependence, whereby the relationship between the extremes of multiple variables evolves over time, is commonly observed in many environmental and financial data sets. However, most multivariate extreme value models are only suited to stationary data. A recent approach to multivariate extreme value modelling uses a geometric framework, whereby extremal dependence features are inferred through the limiting shapes of scaled sample clouds. This framework can capture a wide range of dependence structures, and a variety of inference procedures have been proposed in the stationary setting. In this work, we first extend the geometric framework to the non-stationary setting and outline assumptions to ensure the necessary convergence conditions hold. We then introduce a flexible, semi-parametric modelling framework for obtaining estimates of limit sets in the non-stationary setting. Through rigorous simulation studies, we demonstrate that our proposed framework can capture a wide range of dependence forms and is robust to different model formulations. We illustrate the proposed methods on financial returns data and present several practical uses.

Keywords: Extremal Dependence, Generalised Additive Models, Limit Sets, Multivariate Extremes, Non-stationarity

1 Introduction

1.1 Extremal dependence

Understanding the relationship between the extremes of multiple variables, termed extremal dependence, is an important area of research, with applications in environmental science (Brunner et al., 2016), engineering (Tendijck et al., 2021), actuarial science (Quinn et al., 2019) and finance (Nolde and Zhou, 2021). Typically, modelling procedures aim to estimate fixed coefficients or parameters related to the extremal dependence structure, which are subsequently used to infer extremal statistics. This field of study, known as multivariate

extreme value theory, helps practitioners to carry out robust joint risk assessments across multiple variables, or locations, relevant to a given application.

In the bivariate setting, there exist two distinct subclasses of extremal dependence. Given variables $X_i \sim F_i, i = 1, 2$, define the conditional probability

$$\chi(u) := \Pr(F_2(X_2) > u \mid F_1(X_1) > u) \in [0, 1],$$

and set $\chi = \lim_{u \rightarrow 1^-} \chi(u) \in [0, 1]$, where this limit exists (e.g., Coles et al., 1999). This dependence measure provides a summary of joint tail behaviour, and when $\chi > 0$ and $\chi = 0$, we say that the random vector $\mathbf{X} := (X_1, X_2)$ exhibits asymptotic dependence (AD) and asymptotic independence (AI), respectively. Notably, classical extremal dependence modelling procedures, which are developed under the framework of regular variation, have been shown to perform poorly for data exhibiting AI (Ledford and Tawn, 1997; Heffernan and Tawn, 2004). Since AI is commonly observed in many relevant data sets, this is significant drawback of classical approaches, and many recent works have argued against their use when AI is present (e.g., Opitz, 2016; Murphy-Barltrop et al., 2024b; Huser et al., 2025). In this paper, we pioneer the development of non-stationary extremal dependence modelling via a geometric approach based on limit sets (Nolde, 2014; Nolde and Wadsworth, 2022). The framework can accommodate both AD and AI structures, but the modelling choices we subsequently make are better suited to AI. For simplicity, we also restrict attention to the bivariate setting throughout, noting that many of the introduced techniques can, in principle at least, be extended to higher dimensions; see Section 6 for further discussion.

1.2 Non-stationary extremal dependence

Most modelling frameworks for extremal dependence make the simplifying assumption that data are independent and identically distributed (IID). However, many real-world data sets exhibit non-stationarity, whereby the distribution of the data is not fixed in time. Non-stationarity can be present in two distinct forms: within the marginal distributions F_1, F_2 , and within the (extremal) dependence structure of \mathbf{X} . We restrict attention to the latter form, noting that stationary margins can be obtained by fitting non-stationary marginal models and transforming the data via the probability integral transform to a standard marginal scale. Such transformations do not alter the dependence structure of \mathbf{X} , owing to Sklar’s theorem (Sklar, 1959). However, as recently demonstrated in Kakampakou et al. (2024), selecting the appropriate marginal model for non-stationary data is not straightforward, and poor estimates of the marginal distribution can strongly affect the representation of extremal dependence; see also Murphy-Barltrop and Wadsworth (2024).

In the context of extremal dependence, non-stationarity can result in dependence coefficients or param-

eters that are not fixed in time, requiring modelling procedures that can capture trends in these quantities. Let $\mathbf{X}_t := (X_{1,t}, X_{2,t})$ denote a continuous, non-stationary bivariate process with standardised IID marginal distributions indexed by $t \in \mathbb{Z}$, where t represents discrete-time, and \mathbf{Z}_t is a continuous, p -dimensional covariate process (the predictor variables) that directly influences \mathbf{X}_t (the response vector). Our interest lies in understanding, and modelling, the relationship between the extremal dependence structure of \mathbf{X}_t and the covariate process \mathbf{Z}_t . Relatively few approaches have been proposed in this context, even though non-stationary dependence is a common feature across many relevant data sets (e.g., Jonathan et al., 2014b; Castro-Camilo et al., 2018; Murphy-Barltrap and Wadsworth, 2024). We note here that typically we only observe any given non-stationary process for a discrete number of time points; we therefore let $\{1, \dots, T\}$ denote a finite set of arbitrary, fixed indices.

The majority of the available models for non-stationary dependence are targeted at data structures exhibiting AD (see, e.g., de Carvalho and Davison, 2014; Mhalla et al., 2017; Castro-Camilo et al., 2018; Mhalla et al., 2019), and will therefore perform poorly for data exhibiting AI. In the AI setting, Jonathan et al. (2014b), Guerrero et al. (2023) and Talento et al. (2025) consider non-stationary extensions of the conditional extremes model proposed by Heffernan and Tawn (2004), while Mhalla et al. (2019), Murphy-Barltrap and Wadsworth (2024) and André et al. (2025) propose non-stationary extensions, and inference techniques, for a quantity termed the angular dependence function (Wadsworth and Tawn, 2013). Alongside these approaches, Mhalla et al. (2019) and Lee et al. (2024) model coefficients of extremal dependence in the non-stationary setting; in particular, the limit χ and the coefficient of tail dependence $\eta \in (0, 1]$, first proposed by Ledford and Tawn (1996), which quantifies the dependence amongst asymptotically independent data structures; see Appendix B for a formal definition.

The available approaches for modelling non-stationary extremal dependence offer limited utility. For example, coefficients such as η and χ merely summarise the extremal dependence, and provide no information outside of the region where all variables are jointly large (i.e., the top right corner of the positive quadrant \mathbb{R}_+^2). Furthermore, the Heffernan and Tawn (2004) and Wadsworth and Tawn (2013) modelling frameworks suffer from various drawbacks; the former requires one to fit several separate models, which can result in contradictory conclusions (Liu and Tawn, 2014; Simpson and Tawn, 2024a), while the latter can only be used for specific probability calculations (e.g., joint survivor sets). These limitations, combined with the sparsity of existing approaches, motivate novel developments for modelling non-stationary extremal dependence in the AI setting.

1.3 Limit sets

An increasingly popular modelling tool for multivariate extremes involves the use of limit sets, which arise from the asymptotic convergence of scaled sample clouds for suitably chosen scaling sequences and marginal distributions, to infer extremal dependence properties and perform inference. We briefly summarise this convergence property below. Let $\mathbf{X} = (X_1, X_2)$ be a continuous random vector with a stationary copula and marginal distribution functions asymptotically equal to a von Mises function, i.e., light-tailed margins (Nolde and Wadsworth, 2022). Consider the set $C_n := \{\mathbf{X}_i/r_n\}_{i=1}^n$ denoting n independent copies of \mathbf{X} scaled by a suitably chosen sequence $r_n > 0$ satisfying $r_n \rightarrow \infty$ as $n \rightarrow \infty$. Under mild conditions, C_n converges in probability, as $n \rightarrow \infty$, onto a deterministic set $\mathcal{G} := \{\mathbf{x} : g(\mathbf{x}) \leq 1\}$ (Davis et al., 1988; Kinoshita and Resnick, 1991; Balkema and Nolde, 2010). The set \mathcal{G} is known as the *limit set*, and $g : \mathbb{R}^2 \mapsto \mathbb{R}_+$ is termed the *gauge function* of \mathcal{G} .

The shape of \mathcal{G} is affected by the margins and dependence structure of \mathbf{X} . To focus on dependence only, a common approach is to standardize margins. We adopt standard Laplace margins in this article. The theoretical limit sets for three copulas on Laplace margins are illustrated in Figure 1, alongside the scaled sample clouds C_n for a large, but finite, sample size n . One can observe that even at the finite level, the scaled observations lie approximately within the interior \mathcal{G} . Further background on limit sets and conditions for convergence of C_n are discussed in Section 2.2.




Figure 1: Theoretical limit and boundary sets, denoted as blue shaded regions and thick lines, respectively, for Gaussian (left), student- t (centre) and inverted logistic (right) copulas.

Many recent works have introduced techniques for estimating \mathcal{G} . These approaches can be divided into two categories: those that only target limit set estimation, and those that use the geometric framework as a vehicle for full joint tail inference. For the former, Simpson and Tawn (2024a) and Majumder et al. (2025) introduce semi-parametric techniques using generalised additive models (GAMs) and Bézier splines, respectively, to approximate \mathcal{G} . Note that we also employ GAMs and splines for our proposed modelling approach,

although in an entirely different context (i.e., non-stationary extremal dependence) requiring more complex formulations. For the latter category, Wadsworth and Campbell (2024) propose a parametric approach, Campbell and Wadsworth (2024) present a semi-parametric piecewise linear model, Papastathopoulos et al. (2023) introduce a latent Gaussian model, and Murphy-Barltrop et al. (2024a) put forward a deep learning approach. In each of these works, the authors perform rigorous simulation studies, and demonstrate that the proposed techniques are able to accurately capture a wide variety of extremal dependence structures encompassing both dependence classes.

It is clear that modelling multivariate extremes through limit sets represents an important line of research. Furthermore, existing estimation approaches for \mathcal{G} have clearly demonstrated many advantages of the geometric framework; namely, robustness across different dependence structures (Majumder et al., 2025), self-consistent conclusions (Simpson and Tawn, 2024a), and greater applicability in higher dimensions (up to $d = 10$) than many competing approaches for modelling extremal dependence (de Monte et al., 2025). However, to date, the literature on limit set estimation has exclusively focused on the case when data are IID, and thus cannot be applied for data exhibiting non-stationary dependence.

1.4 Outline

In this article, we introduce a novel non-stationary framework for modelling extremal dependence via limit sets. Our approach can be applied to data exhibiting both AD and AI, with more of a focus placed on the latter class. Moreover, inference from our approach is not limited to summary dependence measures or specific joint tail probabilities: for any index t , our model can be used for simulation from the joint tail of \mathbf{X}_t , for probability estimation in general joint tail regions, and for approximating bivariate risk measures.

In Section 2, we outline a natural extension of the geometric framework to the non-stationary setting, and detail our asymptotic assumptions. In Section 3, we explain our GAM framework for estimating the non-stationary counterpart of the gauge function (equivalently, the limit set). In Section 4 we apply our framework to a range of simulated data sets, demonstrating the proposed methodology can approximate non-stationary boundary sets across many dependence structures. In Section 5, we apply our model to financial data, with diagnostics indicating good performance. We conclude in Section 6 with a discussion and outlook on future work.

2 Limit sets and their non-stationary extensions

2.1 Background on limit sets

The limit set \mathcal{G} of the scaled sample cloud C_n was introduced in Section 1.3. When \mathbf{X} has standard Laplace margins, Balkema and Nolde (2010) and Nolde and Wadsworth (2022) provide a sufficient condition for C_n to converge onto \mathcal{G} .

Proposition 2.1 (Corollary of Proposition 2.2 of Nolde and Wadsworth (2022)). *Suppose the joint density, $f_{\mathbf{X}}(\cdot)$, of \mathbf{X} exists, and that*

$$-\log f_{\mathbf{X}}(u\mathbf{x}) \sim ug(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^2, \quad u \rightarrow \infty, \quad (1)$$

for a continuous gauge function $g(\cdot)$. Then, C_n converges in probability onto \mathcal{G} as $n \rightarrow \infty$.

The limit set \mathcal{G} is star-shaped at $\mathbf{0} := (0, 0)$ and compact (Kinoshita and Resnick, 1991). It is clear from equation (1) that the gauge function $g(\cdot)$ is 1-homogenous, i.e., $g(c\mathbf{x}) = cg(\mathbf{x})$ for any $c > 0$, $\mathbf{x} \in \mathbb{R}^2$, linking to the star-shaped property. In standard Laplace margins, $\mathcal{G} \subseteq [-1, 1]^2$, with the corresponding limit set possessing componentwise maxima and minima $\mathbf{1}$ and $-\mathbf{1}$, respectively. For any random vector exhibiting AD, such as the student- t copula, we have $\mathbf{1} \in \mathcal{G}$ (see Figure 1). Here and in what follows, all convergence results involving sets are understood in the Hausdorff sense; we refer to Nolde and Wadsworth (2022) for a detailed discussion and theoretical treatment of limit sets.

Following Nolde (2014), we also define the *boundary* set $\partial\mathcal{G} := \{\mathbf{x} : g(\mathbf{x}) = 1\} \subset \mathcal{G}$, noting this set possesses the same componentwise maxima and minima as \mathcal{G} . The boundary sets are also illustrated for the three copula examples in Figure 1. We remark that one only needs to know $\partial\mathcal{G}$ in order to specify \mathcal{G} ; hence no knowledge is lost through just considering the boundary set. We therefore use the term ‘limit set’ to refer to both $\partial\mathcal{G}$ and \mathcal{G} throughout this work.

Recent works have demonstrated that the limit set is directly related to extremal dependence structure of \mathbf{X} (Nolde, 2014; Nolde and Wadsworth, 2022). Specifically, $\partial\mathcal{G}$ links several representations for multivariate extremes, immediately providing coefficients and parameters associated with the models proposed in Ledford and Tawn (1996), Heffernan and Tawn (2004), Wadsworth and Tawn (2013) and Simpson et al. (2020). For example, taking $\eta \in (0, 1]$, we have that $\eta := \max\{s \in [0, 1] : [s, \infty)^2 \cap \partial\mathcal{G} \neq \emptyset\}$. In this sense, the study of deterministic geometric shapes is closely linked to the study of extremal dependence, and we refer to this topic as *geometric extremes* herein. Furthermore, in a practical setting, estimates of $\partial\mathcal{G}$ can be used to obtain estimates of risk measures, including, but not limit to, joint tail probabilities (Wadsworth and Campbell, 2024), return curves (Murphy-Barltrop et al., 2024b), and return level sets (Papastathopoulos et al., 2023). Thus, knowledge of $\partial\mathcal{G}$ provides not only a great deal of information about the extremal dependence structure,

but also supplies practitioners with a framework for performing estimation on multivariate extremes.

A brief review of statistical estimation techniques for \mathcal{G} was given in Section 1.3. When performing statistical estimation in the geometric extremes framework, it is helpful to decompose a random vector into angular and radial components. However, the method by which one performs this decomposition is ambiguous, owing to the theoretically infinite number of angular-radial systems that are available; see Mackay and Jonathan (2023) for further discussion. As such, we define angular radial systems in a somewhat general manner, and leave the details of our chosen system(s) to Section 3.

Let $\|\cdot\| : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ denote any valid norm. Given $\mathbf{X} \in \mathbb{R}^2 \setminus \mathbf{0}$, define radial R and angular \mathbf{W} variables by $\mathbf{X} \mapsto (R, \mathbf{W}) := (\|\mathbf{X}\|, \mathbf{X}/\|\mathbf{X}\|)$, such that $R > 0$ and $\mathbf{W} \in \mathcal{S}^1$, where $\mathcal{S}^1 := \{\mathbf{x} \in \mathbb{R}^2 : \|\mathbf{x}\| = 1\}$ denotes the (closed) unit ball of $\|\cdot\|$. Clearly $\mathbf{X} = R\mathbf{W}$, implying the behaviour of \mathbf{X} is directly related to the joint distribution of (R, \mathbf{W}) . It is trivial to show that the mapping $t : \mathbb{R}^2 \setminus \mathbf{0} \mapsto \mathbb{R}_+ \times \mathcal{S}^1$, where $t(\mathbf{x}) := (\|\mathbf{x}\|, \mathbf{x}/\|\mathbf{x}\|)$, is bijective; thus, no information is lost through considering (R, \mathbf{W}) . Owing to the star-shaped property of \mathcal{G} and the homogeneity of $g(\cdot)$, one can then reformulate the boundary set in terms of the angular component; specifically, we have

$$\partial\mathcal{G} = \{r\mathbf{w} : \mathbf{w} \in \mathcal{S}^1, r > 0, g(r\mathbf{w}) = 1\} = \{r\mathbf{w} : \mathbf{w} \in \mathcal{S}^1, r = 1/g(\mathbf{w})\} = \{\mathbf{w}/g(\mathbf{w}) : \mathbf{w} \in \mathcal{S}^1\}.$$

This implies that $\partial\mathcal{G}$ can be specified by only evaluating the gauge function on \mathcal{S}^1 .

All of the existing limit set estimation approaches make model assumptions about the tail of $(R \mid \mathbf{W} = \mathbf{w})$, $\mathbf{w} \in \mathcal{S}^1$, from which the limit set can be approximated. Treating the angle as a covariate, or predictor variable, this modelling strategy is similar to a non-stationary univariate extreme value analysis, which is also the foundation of the recently-developed SPAR model for multivariate extremes (Mackay and Jonathan, 2023), and for which a wide range of modelling approaches are available (see, e.g., Murphy-Barltrop et al., 2025; Mackay et al., 2025).

2.2 Non-stationary limit sets

Let \mathbf{X}_t denote a non-stationary, bivariate process with common margins (e.g., Laplace). To extend the geometric extremes framework to the non-stationary setting, we first require additional notation and assumptions. Letting $\mathbf{z} \in \mathbb{R}^p$ denote any fixed covariate realisation, we make the following assumptions.

Assumption 2.1. *The process \mathbf{X}_t is conditionally stationary w.r.t \mathbf{Z}_t , i.e., $(\mathbf{X}_t \mid \mathbf{Z}_t = \mathbf{z}) \stackrel{d}{=} (\mathbf{X}_{t'} \mid \mathbf{Z}_{t'} = \mathbf{z})$ for all $t, t' \in \{1, \dots, T\}$.*

Assumption 2.2. *The process \mathbf{X}_t is conditionally independent, i.e., $(\mathbf{X}_t \mid \mathbf{Z}_t) \perp\!\!\!\perp (\mathbf{X}_{t'} \mid \mathbf{Z}_{t'})$ for all $t \neq t'$.*

Assumption 2.3. *The joint density of $(\mathbf{X}_t \mid \mathbf{Z}_t = \mathbf{z})$, denoted $f_{\mathbf{X}_t \mid \mathbf{Z}_t}(\cdot \mid \mathbf{z})$, exists and satisfies*

$$-\log f_{\mathbf{X}_t \mid \mathbf{Z}_t}(u\mathbf{x} \mid \mathbf{z}) \sim ug_{\mathbf{z}}(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^2, \quad u \rightarrow \infty,$$

for all $t \in \{1, \dots, T\}$, where $g_{\mathbf{z}}(\cdot)$ is termed the *non-stationary gauge function*.

Let $C_n^{\mathbf{z}} := \{\mathbf{X}'_t/r_n\}_{t=i_1}^{i_n}$, where $\mathbf{X}'_t := \mathbf{X}'_t(\mathbf{z}) = (\mathbf{X}_t \mid \mathbf{Z}_t = \mathbf{z})$ for all $t \in \{i_1, \dots, i_n\} \subseteq \{1, \dots, T\}$ such that $n \leq T$ and $T \rightarrow \infty$ as $n \rightarrow \infty$. Assumptions 2.1–2.3 yield the following result.

Proposition 2.2. *Suppose that $\mathbf{X}_t \mid \mathbf{Z}_t$ obeys Assumptions 2.1–2.3. Then, the scaled sample cloud $C_n^{\mathbf{z}}$ converges in probability onto $\mathcal{G}_{\mathbf{z}} := \{\mathbf{x} : g_{\mathbf{z}}(\mathbf{x}) \leq 1\}$ as $n \rightarrow \infty$, for all $\mathbf{z} \in \mathbb{R}^p$.*

Proof. Under Assumptions 2.1 and 2.2, $\{\mathbf{X}'_t\}_{t=i_1}^{i_n}$ is an IID sequence of n random vectors for all $\mathbf{z} \in \mathbb{R}^p$. Taking Assumption 2.3 and applying Proposition 2.1 pointwisely for all $\mathbf{z} \in \mathbb{R}^p$, the result follows. \square

As in Section 2.1, we define the boundary set $\partial\mathcal{G}_{\mathbf{z}} := \{\mathbf{x} : g_{\mathbf{z}}(\mathbf{x}) = 1\} \subset \mathcal{G}_{\mathbf{z}}$, which we term the *non-stationary limit set* henceforth. Two comments on the connection between theory and practice are in order. First, in practical applications with continuous covariates, one will never observe repeated observations at a fixed realisation \mathbf{z} ; consequently, our estimation strategy, outlined in Section 3, uses an established regression approach to smoothly approximate $\partial\mathcal{G}_{\mathbf{z}}$ over the covariate domain. Second, standard Laplace marginals for \mathbf{X}_t are a natural choice in practice as they offer a more complete picture of extremal dependence compared to single-tailed marginal distributions, especially in the case of negative dependence (Murphy-Barltrop et al., 2025).

We emphasise that the framework outlined in this section is more general than most existing approaches for modelling non-stationary extremal dependence detailed in Section 1.2. In particular, knowledge of $\partial\mathcal{G}_{\mathbf{z}}$ allows us to estimate a wide range of dependence coefficients, parameters and risk measures. For instance, taking the coefficient of tail dependence $\eta \in (0, 1]$, a non-stationary counterpart is given by $\eta_{\mathbf{z}} := \max\{s \in [0, 1] : [s, \infty)^2 \cap \partial\mathcal{G}_{\mathbf{z}} \neq \emptyset\}$. Such quantities are useful for summarising trends in extremal dependence and performing probability estimation. As another example use case, consider the following: given any $p \in (0, 1)$ close to 1, $\partial\mathcal{G}_{\mathbf{z}}$ allows us to specify a set $\mathcal{A}_{\mathbf{z}}^p$ satisfying $\Pr[\mathbf{X}_t \in \mathcal{A}_{\mathbf{z}}^p \mid \mathbf{Z}_t = \mathbf{z}] = p$. Such sets are known as environmental contours (Haver and Kleiven, 2004) or return level sets (Papastathopoulos et al., 2023). In the geometric extremes setting, this set is centred at $\mathbf{0}$ and is computed using high quantiles at a fixed probability level from the conditional variables $(R \mid \mathbf{W} = \mathbf{w})$, $\mathbf{w} \in \mathcal{S}^1$, thus ensuring an equal probability of exceedance in any angular direction. We term $\mathcal{A}_{\mathbf{z}}^p$ a *non-stationary return level set*, and remark that such sets provide interpretable and intuitive summaries of joint extremal risks that are used to aid with design analysis in practice (Mackay and Haselsteiner, 2021); see Section 5 for further discussion.

Finally, the framework outlined in this section can be used for full tail inference on $(\mathbf{X}_t \mid \mathbf{Z}_t = \mathbf{z})$, including, but not limited to, simulation. This represents a significant practical advantage since many applications of multivariate extreme value theory require a means of efficiently simulating synthetic data in the joint tail, e.g., flood risk mitigation (Keef et al., 2013) and catastrophe modelling (Quinn et al., 2019). Moreover, simulated data allows one to estimate non-stationary tail probabilities for $(\mathbf{X}_t \mid \mathbf{Z}_t = \mathbf{z})$; such probabilities are often used to aid with structural engineering (Jonathan et al., 2014a), portfolio management (Nolde and Zhou, 2021) and environmental planning (Gouldby et al., 2017).

3 Modelling and learning from data

Having outlined a non-stationary framework for limit sets in Section 2, we now turn our attention to estimation of the non-stationary gauge function. As illustrated in Section 2.1, we are only required to estimate $g_{\mathbf{z}}(\cdot)$ on any unit ball to obtain the boundary set, and we therefore restrict our estimation domain to \mathcal{S}^1 . Following Section 2.1, we define $R_t := \|\mathbf{X}_t\|$ and $\mathbf{W}_t := \mathbf{X}_t/\|\mathbf{X}_t\|$, and turn our interest to understanding the tail of $(R_t \mid \mathbf{W}_t = \mathbf{w}, \mathbf{Z}_t = \mathbf{z})$ for any given $\mathbf{w} \in \mathcal{S}^1$ and $\mathbf{z} \in \mathbb{R}^p$. Naturally we wish to avoid making strong assumptions about the tail of $(R_t \mid \mathbf{W}_t = \mathbf{w}, \mathbf{Z}_t = \mathbf{z})$ while simultaneously avoiding over-fitting. Furthermore, we desire an estimator that allows us to simultaneously regress over the angular and covariate domains.

3.1 A simplified representation

The framework introduced Section 2.2 is very general, and could represent a large variety of covariate interactions (e.g., continuous and discrete effects, abrupt regime change), making it challenging to specify an framework that can capture all such possible interactions. As such, we opt to simplify the representation for our estimation procedure, and we take two additional steps. First, we define the polar angular variable $\Phi_t := \text{atan2}(X_{2,t}, X_{1,t}) \in [0, 2\pi)$, where atan2 denotes the 2-argument arctangent function. Given any norm $\|\cdot\|$, one can identify any point $\mathbf{w} \in \mathcal{S}^1$ with an angle $\phi \in [0, 2\pi)$: in particular, the transformation $\mathbf{v} : [0, 2\pi) \mapsto \mathcal{S}^1$ given by $\mathbf{v}(\phi) = (\cos(\phi), \sin(\phi))/\|(\cos(\phi), \sin(\phi))\|$ is one-to-one. Consequently, we can represent the unit ball for any norm via the polar angular interval $[0, 2\pi)$, and no information is lost by considering Φ_t in place of \mathbf{W}_t . Furthermore, it is more straightforward to define functions on the interval $[0, 2\pi)$ compared to \mathcal{S}^1 , owing to the fact the former requires only a univariate periodic formulation to be valid. We therefore reformulate the non-stationary gauge function $g_{\mathbf{z}}(\cdot)$ via $m(\phi, \mathbf{z}) := g_{\mathbf{z}}(\mathbf{v}(\phi))$, $\phi \in [0, 2\pi)$; we opt for the notation $m(\phi, \mathbf{z})$ over $m_{\mathbf{z}}(\phi)$ to emphasise the fact that we treat both ϕ and \mathbf{z} as predictor variables in our framework; see Section 3.3. Under this reformulation, we have $\partial\mathcal{G}_{\mathbf{z}} = \{\mathbf{v}(\phi)/m(\phi, \mathbf{z}) : \phi \in [0, 2\pi)\}$;

see Section 2.1.

To further simplify the formulation, we treat the discrete index t as a continuous covariate (i.e., we set $\mathbf{Z}_t = t$ for all t), and explicitly assume that any relevant or unobserved covariates vary smoothly as a function of t . Combined, these steps allow us to simplify the formulation of the framework introduced in Section 2 and henceforth, we use t and Φ_t in place of \mathbf{Z}_t and \mathbf{W}_t , respectively. We stress that our proposed methodology for this simplified representation is specifically adapted for modelling extremal dependence that smoothly varies in time; see also Section 6.

3.2 Modelling the conditional radial tails

To approximate limit sets in the stationary setting, Wadsworth and Campbell (2024) study the tail behaviour of the variable $(R \mid \Phi = \phi)$. While the authors did not consider the non-stationary setting, their derivations extend to this case. In particular, if Assumption 2.3 holds, we have that $f_{R_t \mid \Phi_t}(r \mid \phi) \propto r \exp\{-rm(\phi, t)[1 + o(1)]\}$, $r \rightarrow \infty$, for any $t \in \{1, \dots, T\}$, where $f_{R_t \mid \Phi_t}(\cdot \mid \phi)$ denotes the density function of $(R_t \mid \Phi_t = \phi)$. In addition, their arguments imply that for many dependence structures

$$f_{R_t \mid \Phi_t}(r \mid \phi) \propto r \exp\{-rm(\phi, t)\}[1 + o(1)], \quad r \rightarrow \infty, \quad (2)$$

giving asymptotic equivalence of equation (2) with a gamma kernel. This motivates the following modelling assumption:

$$(R_t \mid \Phi_t = \phi, R_t > r^\tau(\phi, t)) \sim \text{truncGamma}(2, m(\phi, t)), \quad (3)$$

where ‘truncGamma’ is shorthand for the *truncated gamma* distribution with shape and rate parameters 2 and $m(\phi, t)$, respectively, and $r^\tau(\phi, t)$ is the τ -quantile of $(R_t \mid \Phi_t = \phi)$ for all $t \in \{1, \dots, T\}$ and some $\tau \in (0, 1)$ close to 1, i.e., $\Pr(R_t \leq r^\tau(\phi, t) \mid \Phi_t = \phi) = \tau$. Selection of τ is discussed in Section 4. In this context, we refer to observations satisfying $\|\mathbf{X}_t\| > r^\tau(\text{atan2}(X_{2,t}, X_{1,t}), t)$ (i.e., exceeding the radial quantile) as *joint tail* observations.

For our framework, we assume that the assumption outlined in equation (3) holds exactly at some fixed $\tau < 1$ for all t and $\phi \in [0, 2\pi)$. Empirical evidence from many approaches indicates that this assumption is flexible enough to capture a wide range of dependence structures (Wadsworth and Campbell, 2024; Majumder et al., 2025; Murphy-Barltrop et al., 2024a). This resulting inference procedure involves two-steps: estimation of the non-stationary quantile function $r^\tau(\phi, t)$, and estimation of the non-stationary gauge function $m(\phi, t)$. We represent both of these functions using GAMs, an overview of which is provided in Section 3.3.

We remark that one can also leave the gamma shape parameter of equation (3) as a free parameter to

estimate, as has been done in several existing limit set estimation approaches (e.g., Wadsworth and Campbell, 2024; Murphy-Barltrop et al., 2024a). Fixing this quantity to 2 corresponds to assuming that the asymptotic form of equation (2) holds exactly at some finite level. In unreported results, we found that including the shape parameter in the non-stationary framework did not improve the quality of limit set estimates, and in some cases even reduced their quality. We therefore chose to fix the shape at 2, which has the added benefit of reducing parameter variability.

3.3 Modelling covariate interactions with GAMs

GAMs (Wood, 2017) provide a flexible, semi-parametric framework for capturing complex functional forms without requiring rigid modelling assumptions. Typically, GAMs are used in practice to capture the relationship between a response variable and a set of predictor variables; the resulting fitted model can subsequently be used for prediction and interpolation across the joint domain of the predictor variables. Unlike classical regression techniques, which tend to assume simple linear relationships between variables, GAMs exploit smooth functional forms to capture non-linear patterns, alongside complex interactions that may arise between dependent predictor variables. Moreover, in the extreme value context, theoretical properties regarding convergence and asymptotic normality of GAMs have also been established (Yoshida, 2026).

When fitting GAMs in practice, the optimal choice of smoothing functions and tuning parameters is context dependent, and fine tuning is typically required to ensure the model accurately represents the underlying data structure. Consequently, we only describe GAMs in the context of our proposed framework. As noted in Section 3.2, we wish to estimate the functions $r^\tau(\phi, t)$ and $m(\phi, t)$ for all t and $\phi \in [0, 2\pi)$. Thus, under the GAM framework, we view t and ϕ as predictor variables. The interaction between these variables arises from the fact that for a fixed angle ϕ , the conditional radial tail behaviour at ϕ will vary over time, while for a fixed time point t , the extremal dependence structure at t will vary across angles.

We propose the following GAM formulations

$$\log[r^\tau(\phi, t)] = \beta_0 + s_{tp}(t, \phi), \quad \log[m(\phi, t)] = \beta_0^* + s_{tp}^*(t, \phi), \quad (4)$$

where $\beta_0, \beta_0^* \in \mathbb{R}$ are intercept terms and $s_{tp}(\cdot, \cdot), s_{tp}^*(\cdot, \cdot)$ are smooth tensor product splines capturing the interaction between t and ϕ . Tensor splines are formed by taking products of univariate smooth spline functions defined on the predictor variables. Each univariate spline is formed of a piecewise combination of flexible polynomial functions, and the points at which these polynomials are connected are called *knots*. In our case, we use cubic and cyclic cubic splines for t and ϕ , respectively, where latter spline ensures the periodicity over the polar angular variable. Cubic splines possess many desirable properties, such as optimality

in terms of smoothness, high flexibility and computational efficiency (Wood, 2017). More complex modelling frameworks involving adding univariate smooth splines to the tensor splines of equation (4) did not lead to noticeable improvements within our modelling framework. An illustration of tensor product spline basis functions are provided in Appendix A.

Identical GAM formulations are taken for $r^\tau(\phi, t)$ and $m(\phi, t)$; this is due to the fact, as $\tau \rightarrow 1$, these functions are approximately inversely proportional (Wadsworth and Campbell, 2024) and thus will vary in a similar manner over the predictor variable domain(s). Let $\boldsymbol{\beta}$ and $\boldsymbol{\beta}^*$ denote the spline coefficients associated with the tensor product functions $s_{tp}(\cdot, \cdot)$ and $s_{tp}^*(\cdot, \cdot)$, respectively. Denote the number of knots (i.e., the *basis dimensions* of the univariate splines) as κ_t and κ_ϕ for t and ϕ , respectively. Note that the basis dimensions correspond to the flexibility of the model, and their selection typically represents a trade-off; higher dimensions result in greater flexibility but increased parameter variability. In practice, the tensor product functions in equation (4) are also penalised to avoid over-fitting, which commonly occurs for models with large parameter sets. We return to this discussion in Section 4.4. The resulting parameter sets are given by $\boldsymbol{\theta} := (\beta_0, \boldsymbol{\beta})$ and $\boldsymbol{\theta}^* := (\beta_0^*, \boldsymbol{\beta}^*)$, with tuning parameters κ_t and κ_ϕ and joint predictor domain $\mathcal{D} := \{1, \dots, T\} \times [0, 2\pi)$. We now consider estimation of $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^*$.

3.4 Quantile regression for the conditional radial variable

Estimation of $r^\tau(\phi, t)$ represents a quantile regression problem, for which a wide range of approaches are available (Koenker et al., 2017). These techniques are typically based on the pinball loss function, which, when minimised with respect to a parameter set, returns quantile estimates for any probability level τ . Many recent works have incorporated GAM formulations in the context of quantile regression, allowing for flexible semi-parametric inference (e.g., Yu and Moyeed, 2001; Oh et al., 2011; Koenker, 2011; Youngman, 2019). However, there are several issues with these approaches. First, the standard pinball function is piecewise linear and is consequently difficult to optimise, leading to difficulty when incorporating complex model formulations with large parameter vectors, such as (4). Furthermore, some existing approaches require one to manually select ‘smoothing’ parameters (i.e., parameters to mitigate overfitting for highly parametrised models), even though this selection is not straightforward and is crucial for balancing the trade-off between model fit and complexity. Moreover, to estimate GAM parameters, many approaches adopt an asymmetric Laplace model as a working assumption, even though such distributional assumptions are problematic and can result in poorly calibrated, inaccurate quantile estimates (Fasiolo et al., 2021a).

In this work, we employ the approach introduced in Fasiolo et al. (2021a), which overcomes the drawbacks of existing techniques. They introduce a novel technique for fitting GAM-based quantile regression models

and demonstrate that the resulting estimator is well calibrated, stable and efficient. They also introduce a novel loss function which does not suffer from the optimisation issues faced by the pinball function. This novel loss function is penalised to impose smoothness on the resulting splines (i.e., avoid overfitting), resulting in penalty parameters. Penalty parameter estimates are obtained via a Bayesian approach, with a so-called ‘learning rate’ which determines the relative weights of the loss and penalty functions. Selection of basis dimensions is discussed in Section 4.4, and for every t and $\phi \in [0, 2\pi)$, we denote the corresponding quantile estimate by $\hat{r}^\tau(\phi, t)$.

3.5 Restricted maximum likelihood estimation

With quantile estimates obtained over \mathcal{D} , what remains is to estimate the non-stationary gauge function via the model in (3). For this, we employ the framework of Youngman (2019) to estimate the parameter vector $\boldsymbol{\theta}^*$. Taking the log-likelihood function associated with equation (3), this approach employs a technique known as restricted maximum likelihood (REML), whereby the log-likelihood function is penalised to avoid over-fitting, thus allowing for flexible GAM formulations in the likelihood function (Wood, 2011; Wood et al., 2016). In a similar manner to Section 3.4, this results in penalty parameters that also need to be estimated as part of the inference procedure. Under the restricted maximum likelihood framework, this is achieved by making distributional assumptions on the smoothing parameters, then maximising the resulting likelihoods. For rigorous mathematical details and guidance on model fitting, see Wood (2017) and Youngman (2019).

Given t and $\phi \in [0, 2\pi)$, we denote the corresponding non-stationary gauge function estimate by $\hat{m}(\phi, t)$, with the corresponding non-stationary boundary set estimate given by $\widehat{\partial\mathcal{G}}_t = \{\mathbf{v}(\phi)/\hat{m}(\phi, t) : \phi \in [0, 2\pi)\}$. Owing to the continuity of $s_{tp}^*(\cdot, \cdot)$ (equation (4)), the limit set associated with this estimate is always compact (Murphy-Barltrop et al., 2024a). However, there is no guarantee that $\widehat{\partial\mathcal{G}}_t \in [-1, 1]^2$, or that the boundary set estimate possesses componentwise maxima and minima $\mathbf{1}$ and $-\mathbf{1}$, respectively. See Sections 5 and 6 for further discussion.

Thus far, we have provided a general framework for inference on non-stationary boundary sets, but we have been intentionally ambiguous regarding specific features of the model. In particular, we have not yet specified a norm $\|\cdot\|$, the quantile level τ , or the basis dimensions κ_t and κ_ϕ ; these are important quantities for model inference, and as it turns out, the appropriate values vary depending on the form(s) of extremal dependence. Furthermore, one must also consider the effect of the observation window T and account for variability for any given sample size. Selection of these features is discussed in Sections 4 and 5.

3.6 Model checking

To assess model fits, we introduce two diagnostic tools. Example applications of each diagnostic discussed in this section can be found in Section 5. Since we do not observe repeated observations for any time point t , one cannot compute scaled sample clouds and compare these to fixed boundary set estimates (see, e.g., Figure 1), as is common for many stationary approaches. Therefore, a more nuanced approach is required for assessing model fits in the non-stationary setting.

Firstly, we adopt the goodness-of-fit metrics proposed by Wadsworth and Campbell (2024) and Campbell and Wadsworth (2024). In particular, we compute probability-probability (PP) and quantile-quantile (QQ) plots associated with the truncated gamma model fit. For the former, we produce two plots, one comparing observed and theoretical probabilities, and the other illustrating the differences between these probabilities over the observation rank. For the latter, all threshold exceeding observations are transformed to a standard exponential scale using equation (3). These transformed observations can be compared to theoretical quantiles. Combined, these metrics provide global diagnostics for model fits over both the polar angle and time variables. Confidence bounds can be added to these plots by exploiting the fact the k -th order statistic of a uniform distribution theoretically follows a Beta($k, n + 1 - k$) distribution (David and Nagaraja, 2004).

Next, we consider the return level set diagnostic proposed in Murphy-Barltrop et al. (2024a). Introduced in the stationary setting, this model computes the empirical probability of points inside an estimated return level set $\hat{\mathcal{A}}_t^p$ for some p close to one, and compares this value to the nominal level. Repeating this procedure over a range of probability levels, one obtains a QQ plot representing the nominal and estimated probabilities. This diagnostic is easily extended to the non-stationary setting by averaging return level set probabilities over time. Assuming unbiased estimation of return level sets, we have that

$$\frac{1}{T} \sum_{t=1}^T \Pr(\mathbf{X}_t \in \mathcal{A}_t^p) = \sum_{t=1}^T \frac{p}{T} = p. \tag{5}$$

Equation (5) can be approximated empirically via the technique described in Murphy-Barltrop et al. (2024a). In particular, given observations \mathbf{x}_t , $t = 1, \dots, T$ with T large, we set $\hat{p} = (1/T) \sum_{t=1}^T \mathbb{1}(\|\mathbf{x}_t\| \leq \hat{r}^p(\text{atan2}(x_{2,t}, x_{1,t}), t))$, where $\mathbb{1}(\cdot)$ is the indicator function and $\hat{r}^p(\cdot, \cdot)$ denotes the estimated p -quantile function of $(R_t \mid \Phi_t = \phi)$ (i.e., the radial values of $\hat{\mathcal{A}}_t^p$) derived from the truncated gamma model fit. The corresponding diagnostic describes how well the fitted model captures the observed structure in the data, and indicates the suitability of the model for inferring risk measures.

4 Simulation study

In this section, we explore the performance of the framework introduced in Sections 2 and 3. Section 4.1 introduces a range of simulated examples exhibiting non-stationary dependence. In Section 4.2, we discuss metrics for assessing model performance over \mathcal{D} . Section 4.3 discusses estimation of the penalty parameters from the frameworks introduced in Sections 3.4 and 3.5. Model formulation is discussed in Section 4.4 and in Section 4.5, we present results for each of the simulated examples. To estimate the quantile and gauge functions, we employ the `qgam` (Fasiolo et al., 2021b) and `evgam` (Youngman, 2022) software packages, respectively, in the R programming language. Example code for fitting our framework is available at <https://github.com/callumbarltrop/NSGE>, and wrapper functions are included into the GitHub version of `evgam`; see <https://github.com/byoungman/evgam>.

4.1 Simulated examples of non-stationary dependence structures

Given an observation window T , we define a range of non-stationary structures in terms of the time covariate t . All samples are simulated on standard Laplace margins. The first two examples are obtained using the bivariate Gaussian copula, for which dependence is characterised by the correlation coefficient $\rho \in [-1, 1]$. For the first case, we set $\rho(t) := 0.2 + 0.6(t-1)/(T-1)$, giving $\rho(1) = 0.2$ and $\rho(T) = 0.8$, i.e., moving from weak to strong positive dependence. For the second example, we define the functions $a(t) := 0.45(t-1)/(T-1)$ and $b(t) := 0.5 \sin(2.5\pi(t-1)/(T-1))$, and set $\rho_2(t) = a(t) + b(t)$; this corresponds to a harmonic, increasing correlation function. The dependence trend is complex in this case, moving from positive, to negative, then back to positive dependence. An illustrative figure of $\rho_2(t)$ over t is provided in Appendix A.

For the third example, we use the inverted bivariate extreme value copula with the logistic family (Ledford and Tawn, 1997). Dependence is classified through the parameter $\alpha \in (0, 1)$, with positive dependence increasing as α approaches 0. We set $\alpha(t) := 0.3 + 0.4(t-1)/(T-1)$, resulting in a trend that moves from strong to weak positive dependence.

For the fourth example, we consider the student- t copula family, for which dependence is quantified via a parameter $\sigma \in [-1, 1]$ representing the degree of collinearity in the data, alongside the degrees of freedom parameter $\nu > 0$. Lower values of ν correspond with stronger tail dependence. In this case, we fix $\sigma = 0.5$ for all t and set $\nu(t) := 0.5 + 1.5(t-1)/(T-1)$; this corresponds to a trend whereby the collinearity parameter remains constant over time, but the tail dependence gradually decreases.

Finally, for the fifth example, we consider the copula model introduced by Huser and Wadsworth (2019). This model was first proposed on one-tailed margins, and consequently we reformulate it slightly to define it with standard Laplace margins. A version of this model with asymmetric tails has also been considered by

Gong and Huser (2022). Let S denote a standard Laplace variable and \mathbf{V} denote a bivariate random vector with Gaussian copula on standard Laplace margins with correlation coefficient ρ , with S and \mathbf{V} independent from each other. Given any $\delta > 0$, define the random vector

$$\mathbf{Y} = \begin{cases} \delta(S, S) + \mathbf{V}, & \delta < 1, \\ (S, S) + \mathbf{V}/\delta, & \delta \geq 1, \end{cases} \quad (6)$$

where addition is applied componentwise. Note that \mathbf{Y} does not possess standard Laplace margins, and additional marginal transformation involving the parameter δ is required to obtain standardised margins. The model in equation (6) can be interpreted as follows: for $\delta < 1$, \mathbf{V} is heavier tailed than δS , inducing AI, while for $\delta \geq 1$, S dominates \mathbf{V}/δ , resulting in AD. One can observe that a Gaussian copula is recovered as we let δ approach 0. The transition between the two dependence classes occurs at $\delta = 1$, corresponding to an interior point of the parameter space; this is in contrast to most extreme value models, for which this transition occurs at the boundary of the parameter space(s) (Wadsworth et al., 2017). Consequently, we fix $\rho = 0.5$ and set $\delta(t) := 0.7 + 1.8(t - 1)/(T - 1)$, corresponding to a smooth transition between AI and AD.

We simulate 200 samples from each of the copula structures outlined in this section. This allows us to assess both the bias and variability of our proposed modelling framework across the different dependence structures. For each copula example, illustrations of the resulting boundary sets over time are provided in Figure 2. We note that the cases where ‘pointiness’ is observed in some, or all, of the quadrant corners (e.g., the student- t) correspond to copulas possessing strong tail dependence within such regions.

4.2 Performance metrics

The overall goal of our modelling framework is to estimate non-stationary boundary sets, and consequently we assess performance in terms of set estimates. However, owing to the additional time dimension, assessing performance is not straightforward in the non-stationary setting. For example, one could plot all two-dimensional boundary set estimates on a time axes alongside the true sets, but the corresponding three-dimensional plot would be difficult to interpret, and it may be unclear in which regions the model performs well. Furthermore, many performance metrics proposed in the stationary setting, such as mean integrated squared error of the gauge function over $[0, 2\pi)$ (Murphy-Barltrop et al., 2025) or bias in extremal coefficients (Simpson and Tawn, 2024a; Majumder et al., 2025), cannot be easily generalised to the non-stationary setting.

Consequently, we follow Murphy-Barltrop and Wadsworth (2024), who propose fixing either the time index or the polar angle, then evaluating performance visually over the domain of the other component. Firstly, given a set of fixed time points $\{1, \lfloor T/2 \rfloor, T\}$, corresponding to the start, middle and end of the

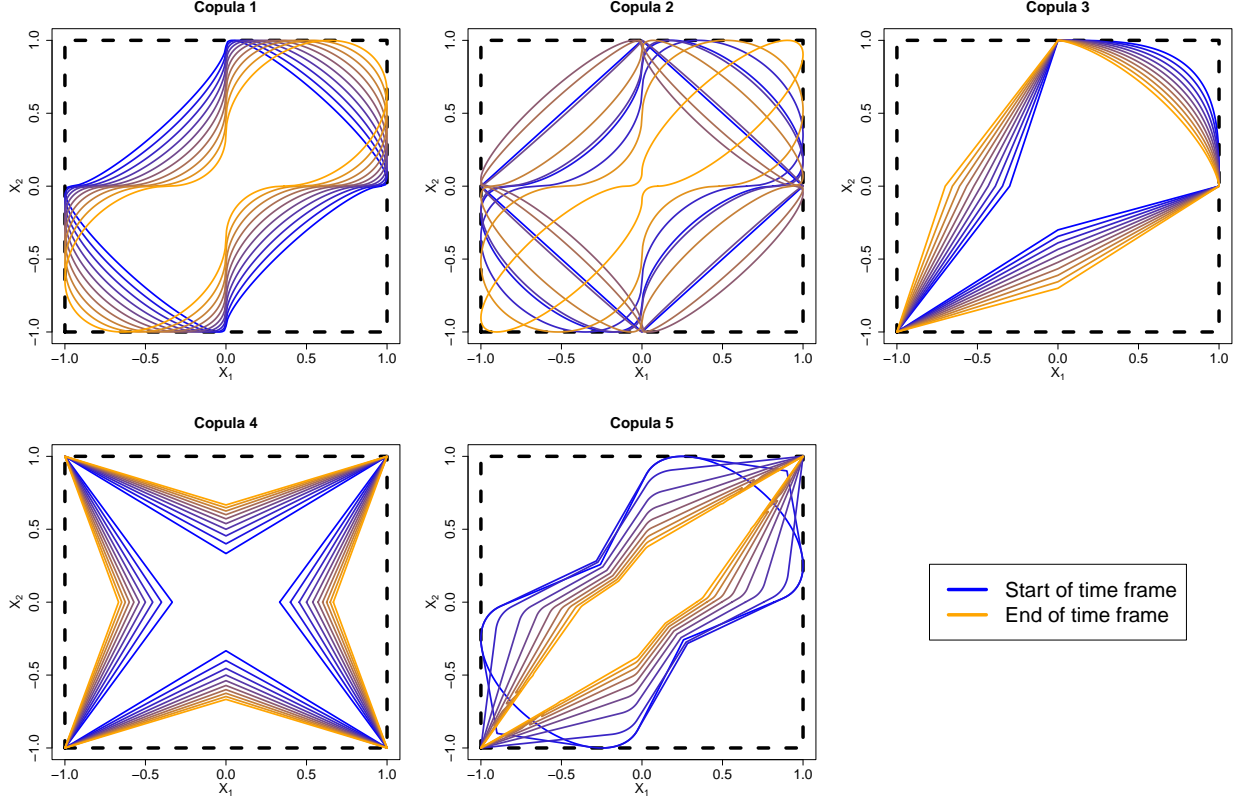


Figure 2: Boundary sets for each copula example at a finite, increasing set of equally spaced time indices. The boundary sets at the start and end of the observation interval are illustrated in blue and orange, respectively, with a colour transition used to visualise the trend in dependence.

observation window, respectively, we compute estimates of the non-stationary boundary set at each $t \in \{1, \lfloor T/2 \rfloor, T\}$ and plot these alongside the true boundary sets. This allows us to visually assess what boundary set shapes can be represented by the framework proposed in Section 3. Evaluating the joint tail at the boundaries of the observation window also provides a good indicator of the overall model performance, since there is less data available for regression at these points. Secondly, we fix angles $\{\pi/4, 3\pi/4, 5\pi/4, 7\pi/4\}$, corresponding to the main diagonals in all four quadrants. For each $\phi \in \{\pi/4, 3\pi/4, 5\pi/4, 7\pi/4\}$, we estimate the corresponding Cartesian point on the boundary set, denoted $\hat{\mathbf{x}}_t^\phi := \mathbf{v}(\phi)/\hat{m}(\phi, t) \in \widehat{\partial\mathcal{G}}_t$, for every t . We then compute the Euclidean distance of $\hat{\mathbf{x}}_t^\phi$, i.e., $\hat{r}_t^\phi := \sqrt{(\hat{x}_{1,t}^\phi)^2 + (\hat{x}_{2,t}^\phi)^2}$, and plot this estimate against t , alongside the corresponding true values. This allows us to visually assess how well our framework is capturing the simulated dependence trends in different regions. Note that we evaluate the Euclidean distance of the Cartesian point to allow for comparison between different norm definitions of the radial component; plotting the estimates of $\hat{m}(\phi, t)$ for different norm definitions, for instance, would be meaningless, since the coordinates at which the gauge function is evaluated are dependent on the choice of norm.

4.3 Selecting GAM smoothing parameters and learning rates

We note that both steps of our modelling procedure, as outlined in Sections 3.4 and 3.5, require smoothing parameters for penalising loss functions used for optimisation. Such parameters can be estimated automatically by making distributional assumptions directly on the smoothing parameters (Wood, 2003). Moreover, for the Bayesian quantile regression framework, a learning rate parameter must be estimated prior to sampling from the posteriors of the GAM parameters. Estimation of, or obtaining samples from, the smoothing parameters is computationally expensive. Moreover, across different data sets from the same copula and sample size, it is unlikely the same smoothing parameters will be estimated twice, implying one cannot easily compare GAM coefficient estimates for different copula samples. Therefore, prior to running our large simulation study, we opt to fix model formulations to speed up computation and ensure comparability across sampling iterations.

For each of the possible combinations of model tuning parameters and features (see Section 3.5), we fit our modelling framework to 10 unique samples from each copula and use the procedures outlined in Fasiolo et al. (2021a) and Youngman (2019) to automatically select the relevant parameters. For the Bayesian quantile regression procedure, we compute the median learning rate across the 10 samples and then fix the learning rate to this estimate for the remainder of the simulation study. Note that being in Bayesian framework, this still requires one to sample from the posteriors of the GAM parameter distributions, and there is no guarantee such distributions will be the same across samples. For the frequentist REML framework, we compute the median smoothing parameters for the tensor product spline and then fix the smoothing parameter for any given combination to these medians. This ensures the same penalised likelihood function is used for each setup.

4.4 Selecting a model formulation

As noted in Section 3.5, our proposed framework requires us to specify a range of tuning parameters and features. Trying to select all of these terms simultaneously would be very difficult, and it is unlikely that the same model formulation would be optimal across all simulated examples. Consequently, we select each model component separately, loosely assuming the optimal choices of each model feature can be selected independently of others.

For the norm function $\|\cdot\|$, we consider three popular choices: $\|\mathbf{x}\|_1 = |x_1| + |x_2|$, $\|\mathbf{x}\|_2 = \sqrt{x_1^2 + x_2^2}$ and $\|\mathbf{x}\|_\infty = \max\{|x_1|, |x_2|\}$, corresponding to the L^1 , L^2 and L^∞ norms, respectively. As observed in Murphy-Bartrop et al. (2025), the unit ball associated with each norm directly influences the shape of the estimated boundary set. To see this, consider Figure 3: in this plot, we have defined a cyclic cubic spline in polar

coordinates, and for each norm, the corresponding Cartesian coordinates have been computed. While the spline is smooth on the polar scale, the resulting shapes in Cartesian coordinates possess noticeable kinks for the L^1 and L^∞ norms. This occurs due to the shapes of corresponding unit balls, which are also illustrated in right panel of Figure 3. Consequently, certain norm choices may be more appropriate for scaled sample clouds which appear particularly ‘pointy’ within certain regions. We stress here that in practice, it is unlikely the same spline would be estimated across all three norms, and Figure 3 is meant to merely illustrate the difference in shapes when splines are transformed to the Cartesian scale.




Figure 3: Plot illustrating the affect of norm choice on Cartesian set constructions. The left and centre panels illustrate the spline functions on polar and Cartesian scales, respectively. The right panel illustrates the unit balls for each choice of norm.

For the quantile level, we consider $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$. This choice represents a typical bias-variance trade-off that is often observed in extreme value theory (see, e.g., Coles, 2001; de Carvalho et al., 2026), and for the framework proposed in Section 3, this equates to assessing the rate of convergence of equation (3) to the true asymptotic tail distribution of the conditional radii variables. We wish to select τ as high as possible without compromising the robustness or reliability of the modelling framework.

For the time basis dimension, we consider $\kappa_t \in \{5, 10, 15\}$ and define spline knots at equally spaced points along the observation window, i.e., for $j = 1, \dots, \kappa_t$, set $k_j^t := 1 + (T - 1)(j - 1)/(\kappa_t - 1)$. As noted in Wood (2017), the exact location of spline knots is not so important, so long as the basis dimension is sufficiently high enough to provide adequate flexibility. For the polar angular basis dimension, we consider $\kappa_\phi \in \{9, 17\}$ and again define knots at equally spaced points on the angular window, i.e., for $j = 1, \dots, \kappa_\phi$, set $k_j^\phi := 2\pi(j - 1)/(\kappa_\phi - 1)$. These choices are practically motivated by the fact such basis dimensions result in knots that include the axes and principal diagonals in \mathbb{R}^2 ; these regions are often where the boundary set exhibits key features related to extremal dependence properties. As such, it is important that our modelling framework can capture the behaviour around these regions. An illustrative figure of the polar angular knot

locations is provided in Appendix A, and we refer to Simpson and Tawn (2024b) for further discussion.

To begin, we consider the effect of the quantile level τ . For this, we first fix the basis dimensions to their maximal values, i.e., $\kappa_t = 15$, $\kappa_\phi = 17$, so as to provide maximal flexibility, and just consider the L^2 norm. For each example in Section 4.1, we also set $T = 25,000$ to provide a large, but not unreasonable, sample size. Note that smaller T values are subsequently considered once all tuning parameters have been selected. For each copula example, we then apply the framework introduced in Section 3 across 200 examples and compute the performance metrics discussed in Section 4.2. These results are given in Appendix C.1. Encouragingly, the model appears remarkably robust to the choice of τ , accurately capturing the range of dependence structures over \mathcal{D} . Naturally there is higher variability at higher τ levels, as one would expect, though we also observe slightly better coverage for such values. Consequently, we fix $\tau = 0.8$ for the remainder of this article; this value appears to offer a reasonable trade-off between bias and variance.

Next, we assess the effect of the basis dimensions κ_t and κ_ϕ . For the former, we fix $\kappa_\phi = 17$ and again restrict attention to the L^2 norm, then vary $\kappa_t \in \{5, 10, 15\}$. Using the performance metrics described in Section 4.2, the results are given in Appendix C.2. For most of the copula examples, there appears to be little difference among the basis dimensions. However, for the complex dependence trend exhibited by copula 2, $\kappa_t = 5$ does not appear sufficiently flexible to capture the underlying trends. Maintaining the same setup with $\kappa_t = 15$ and allowing κ_ϕ to vary in $\{9, 17\}$, we evaluate the effect of basis dimension for the angular component, with the results again given in Appendix C.2. It is clear that the additional flexibility arising from setting $\kappa_\phi = 17$ allows the framework to better capture the complex variety of boundary set shapes. Taking these combined results into account, we fix $\kappa_t = 10$ and $\kappa_\phi = 17$ for the remainder of the article.

Finally, we consider the impact of the norm choice on the modelling framework. For this, we restrict attention to the first and fourth copula examples, where the underlying boundary sets are smooth and pointy respectively. With all other components of the model specified, we vary the norm across the three choices (i.e., L^1 , L^2 and L^∞). These results are given in Appendix C.3. It is clear that the choice of norm does not appear to significantly alter the rate of convergence for the modelling framework described in equation (3), but unsurprisingly affects the types of Cartesian shapes that can be represented by the framework. Visually, it appears that the L^2 and L^∞ norms would be preferable for the first and fourth examples, respectively. However, without knowledge of theoretical boundary sets, it is not clear how one should select the norm choice in practice; we refer to Section 6 for further discussion. Consequently, for the remainder of this simulation study, we restrict attention to the L^2 norm, acknowledging that this will impose smoothness on the resulting boundary set estimates.

4.5 Results

Using the selected tuning parameters and model features, we present our simulation results. Fixing $T = 25,000$, Figure 4 illustrates the boundary set estimates, along with the theoretical boundary sets, for $t = 1$ and $t = T$ with the second, fourth and fifth copula examples; the results for the remaining copulas and time points are given in Appendix C.4. As can be observed, the estimated boundary sets capture the structures of the theoretical sets at the start and end of the time frame.

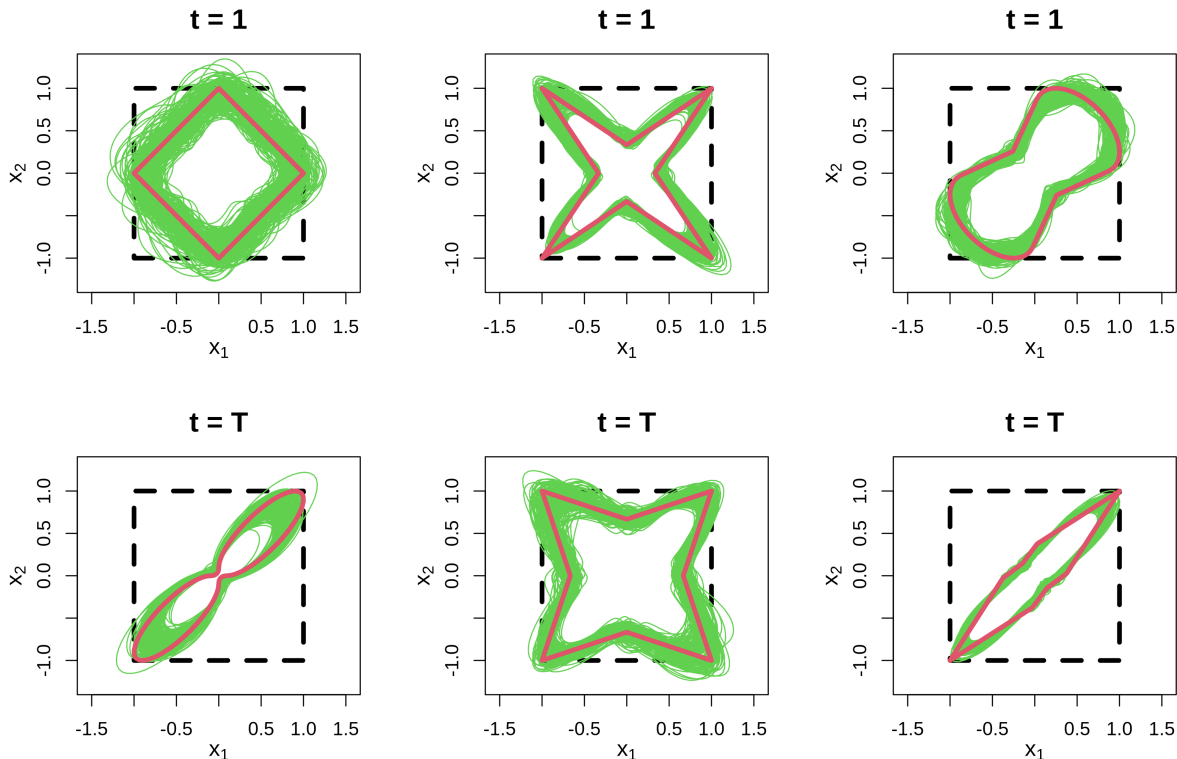


Figure 4: Boundary set estimates for $t = 1$ (top row) and $t = T$ (bottom row). The left, centre and right columns correspond to the second, fourth and fifth copula examples, respectively. The green lines correspond to the resampled set estimates, while the red solid and black dotted lines denote the true boundary set and the region $[-1, 1]^2$, respectively.

Setting $\phi = \pi/4$ and $\phi = 3\pi/4$ and again considering the second, fourth and fifth copula examples, the estimated radii of the corresponding boundary set points over time are illustrated against the truth in Figure 5, with the figures for the remaining angles and copulas given in Appendix C.4. The framework appears to capture the dependence trends in most cases, especially for the Gaussian copulas. However, the model can fail to capture specific dependence features: for example, the abrupt transition between AI and AD for the copula described by equation (6), and the ‘pointiness’ of the student- t copula in the corners. Given the limited amount of knowledge available generally within the non-stationary setting, estimating

dependence features exactly is seldom possible, and we already know that the L^2 norm will not be able to represent the pointedness that arises in certain quadrants. On the whole, however, the slight bias in radii observed at certain angles is secondary to the fact that the estimated Cartesian sets provide reasonable approximations of the true limit sets for each of the copula examples.

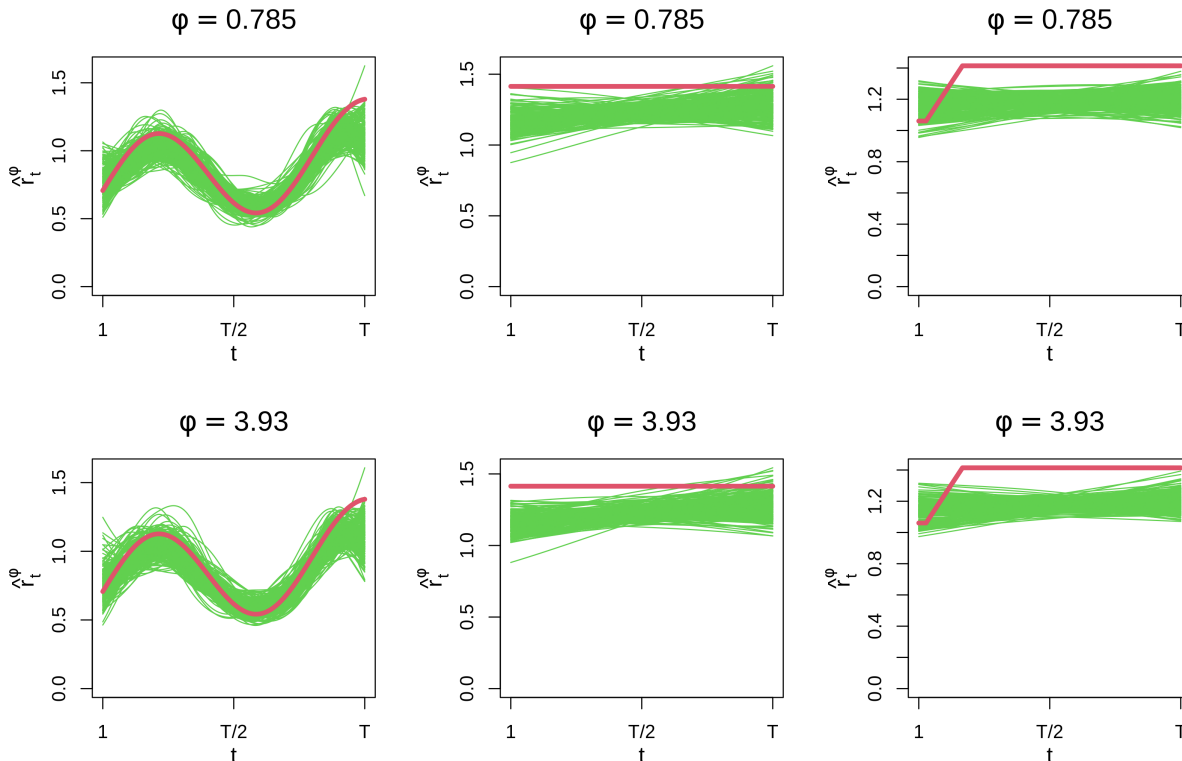


Figure 5: Boundary set radii estimates over time for $\phi = \pi/4$ (top row) and $\phi = 3\pi/4$ (bottom row). The left, centre and right columns correspond to the second, fourth and fifth copula examples, respectively. The legend is identical to that of Figure 4.

Finally, we vary the observation period $T \in \{5,000, 10,000, 25,000\}$ and assess the corresponding boundary set estimates. These results are given in Appendix C.5. Unsurprisingly, increasing the sample size improves bias and reduces variability. However, even for $T = 5,000$, the model still appears able to approximate the shapes of the true boundary sets, and thus is somewhat robust to lower sample sizes.

Overall, our proposed framework is able to accurately capture a wide range of extremal dependence trends, demonstrating the robustness and versatility of the model. Furthermore, the choices of tuning parameters and features generally appear to have little overall influence on the model fits, provided an adequate amount of flexibility is permitted.

5 Case study

5.1 Financial context and data description

We now apply our modelling framework to the NASDAQ stock market. Specifically, we consider data from Apple, Microsoft, Amazon, Google, and Nvidia; these firms are all part of the ‘Magnificent Seven’ (M7), a term which refers to a group of seven highly influential big tech companies. As of September 9th, 2025, these firms constitute the five largest companies (in terms of market capitalisation) in the world. We do not consider the remaining M7 companies (Meta and Tesla) due to limited observational periods of the corresponding time series. In recent years, M7 companies have significantly impacted the U.S. stock market, driving a substantial portion of its growth. Their dominance stems from their leadership in key sectors like consumer electronics, cloud computing, social media, artificial intelligence, and electric vehicles. See Peprah et al. (2025) and de Carvalho and Palacios Ramirez (2025) for further discussion on M7 and on extreme losses of ‘big tech’ stocks, respectively. Owing to the fact these companies began publicly trading on different days, the observational dates, and hence sample size, vary depending on which stocks are being considered. The listing days range from December 12th, 1980 (Apple) to August 19th, 2004 (Google). Daily prices data was accessed from NASDAQ on August 25th, 2025, resulting in time series ending on August 22nd, 2025.

For each company, we take the daily opening price P_t and compute *log-returns* $Q_t := \log(P_t/P_{t-1})$, for all $t > 1$. This quantity represents a mathematically convenient way of representing price changes in financial markets (Tsay, 2010). To avoid computational issues, we remove any time indices t for which $Q_t = 0$ for any stock; this corresponds to two-day periods where there was no change in the opening price of at least one stock. Including such values would result in mixed discrete-continuous time series, which would violate the assumptions of our modelling framework.

In the context of financial risk management, many stakeholders are interested in understanding the joint tails of log-returns across different assets. For example, one may wish to analyse whether an extreme increase or decrease in one stock price will also result in similar changes across other stocks. This motivates the use of multivariate extreme value theory to evaluate joint tail behaviour in regions where variables are simultaneously extreme—either in the lower or upper tails. Considering the bivariate setting, this context represents an application for which tail behaviour across all quadrants is of interest; hence, statistical techniques restricted to a single quadrant (i.e., where both variables are simultaneously large) offer less utility in this setting.

In Section 5.2, we describe our procedure for modelling the marginal distributions of each time series in order to transform to standard Laplace margins. We apply our modelling framework in Section 5.3; the diagnostics subsequently presented in Section 5.4 indicate reasonable model fits across all pairs. Finally, in

Section 5.5, we summarise dependence trends over the observation period and introduce further practical use cases for our proposed framework.

5.2 Marginal modelling of returns

To apply the framework introduced in Section 2.2, we first require marginally IID observations on standard Laplace margins. This requires careful consideration for the log-returns data in question. In particular, log-returns series are unlikely to be independent; see McNeil et al. (2015) for detailed discussion. By computing the autocorrelation function (ACF) of Q_t and Q_t^2 at different lags for each stock, it is clear that our data exhibit non-negligible temporal dependence; see Appendix D for the corresponding plots.

To account for the observed temporal dependence, we model the log-returns using a GARCH(1,1) process and compute the standardised residuals, which we then take to be independent. For this, we assume $Q_t = \mu + \sigma_t \varepsilon_t$, where $\mu \in \mathbb{R}$ denotes a mean term, ε_t denotes the residual process (with zero mean and unit variance), and the conditional volatility σ_t^2 satisfies

$$\sigma_t^2 = c + \alpha Q_{t-1}^2 + \beta \sigma_{t-1}^2,$$

with parameters $c > 0$, $\alpha, \beta \geq 0$ and $\alpha + \beta < 1$. After fitting the model, the standardised residuals $\hat{\varepsilon}_t = (Q_t - \hat{\mu}) / \hat{\sigma}_t$, which we term the *filtered log-returns*, should be approximately independent and identically distributed if the GARCH model is well specified. This works because the GARCH filter removes serial dependence in the conditional variance, capturing volatility clustering and leaving residuals that behave more like white noise (Engle, 1982; Bollerslev, 1986).

The plotted time series of the filtered log-returns are given in Appendix D, alongside ACF plots for $\hat{\varepsilon}_t$ and $\hat{\varepsilon}_t^2$. The latter plots indicate that the standardised residuals are independent in time. Furthermore, the plotted series do not seem to suggest any obvious periodicity or long term trends are present, suggesting the data may exhibit marginal stationarity. We additionally tested this property by applying Augmented Dickey–Fuller (Dickey and Fuller, 1979) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS, Kwiatkowski et al., 1992) tests to each filtered time series. In almost all cases, both tests at a 5% significance level indicated a stationarity assumption was reasonable for the data; the only exception was the KPSS test on the Apple time series for the Apple–Google pairing. Given the same test indicated stationarity for Apple series in all other pairings (for which the observation dates in Apple–Google pairing are included), and the computed p -value was just below the $\alpha = 0.05$ significance level, it is therefore reasonable to assume that this is a type 1 error that can be ignored. We henceforth assume that the filtered log-returns are marginally IID.

As an additional check, we compute empirical bivariate extremograms via the `extremogram` R package

(Frolova and Cribben, 2016) for all pairs of squared residuals. These plots are illustrated in Appendix D and indicate vanishing joint tail dependence, suggesting one can assume that extremes between filtered series are approximately independent across different time lags. We refer to Davis and Mikosch (2009) for further discussion.

We now employ standard techniques to model the individual marginal distributions. In particular, we employ the semi-parametric approach of Coles and Tawn (1991), where a generalised Pareto distribution (GPD) is fitted to the tails while the body is modelled empirically. In the case of returns data, we are interested in understanding both the extreme high and low values. Letting \mathcal{E} denote a filtered log-returns series with observations $\{\varepsilon_1, \dots, \varepsilon_T\}$ and given some α close to 0, define l and h to be the empirical α and $(1 - \alpha)$ quantiles, respectively, of \mathcal{E} . We use the following marginal model

$$\hat{F}_{\mathcal{E}}(\varepsilon) = \begin{cases} 1 - \alpha \{1 + \xi_1(\varepsilon - h)/\sigma_1\}_+^{-1/\xi_1}, & \text{for } \varepsilon > h, \\ \tilde{F}_{\mathcal{E}}(\varepsilon), & \text{for } l \leq \varepsilon \leq h, \\ \alpha \{1 + \xi_2(l - \varepsilon)/\sigma_2\}_+^{-1/\xi_2}, & \text{for } \varepsilon < l, \end{cases}$$

where (σ_1, ξ_1) and (σ_2, ξ_2) are the estimated GPD parameters from the conditional tail variables $(\mathcal{E} - h) \mid \mathcal{E} > h$ and $(l - \mathcal{E}) \mid l > \mathcal{E}$, respectively, $\{x\}_+ := \max\{0, x\}$, and $\tilde{F}_{\mathcal{E}}$ is an empirical rank transform given by $\tilde{F}_{\mathcal{E}}(\varepsilon) = \sum_{t=1}^T \mathbb{1}(\varepsilon_t \leq \varepsilon)/(T + 1)$.

The selection of α is equivalent to selecting a threshold for the GPD, which represents a bias-variance trade-off in practice. For this, we initially applied a leading automated threshold selection technique (Murphy et al., 2025) which outperforms many existing approaches. However, we found the model fits from the Murphy et al. (2025) approach to be inadequate for several time series. Therefore, we instead manually tested a range of α values and eventually fixed $\alpha = 0.03$; this appeared to give reasonable model fits for all of the tails. The resulting GPD QQ plots are given in Appendix D.

With marginal models specified, the probability integral transform is applied to obtain each series on standard Laplace margins. As an additional check, the series are split into non-overlapping rolling windows, and maximum likelihood techniques are used to estimate the location and scale parameter of the Laplace distribution for each window. These rolling window estimates remain close to the true Laplace parameters (i.e., zero location, unit scale) in all cases, indicating the margins have been successfully standardised. These plots are given in Appendix D.

5.3 Modelling time-varying extremal dependence of returns

Time-varying joint modelling in stock markets has been recognised as fundamental by many works (Poon et al., 2003; Patton, 2006; Salvatierra and Patton, 2015; Castro-Camilo et al., 2018), and we now use the framework proposed in Section 3 to this end. Since our model is bivariate, there are a total of 10 pairs of variables we could consider. For the sake of brevity, we restrict our analysis to consider the extremal dependence between Apple and all of the remaining stocks. Apple has been listed for longer than any of the remaining companies, thus representing the most ‘mature’ company, and the stock movements of Apple are often seen as an indicator of broader market trends, particularly within the technology sector. The size of observation windows varies from $T = 6,598$ for the Apple–Nvidia pair to $T = 9,432$ for Apple–Microsoft. As demonstrated in Section 4, such sample sizes are sufficient for applying the modelling framework proposed in Section 3.

Figure 6 illustrates scatterplots of each bivariate data set on Laplace margins, processed as described in Section 5.2. We just show observations in the joint tail (i.e., observations exceeding the estimated threshold function). A colour gradient is used to illustrate the change in joint behaviour over time, suggesting subtle variations in dependence. Ideally, one would wish to formally test for this feature; however, there are no best practices established for such testing within the literature, and we found a previous technique, which involves computing dependence coefficients across overlapping rolling windows (Castro-Camilo et al., 2018; Murphy-Barltrop and Wadsworth, 2024), to be unreliable and misleading. In particular, for simulated IID data, the rolling windows approach could suggest spurious trends when using point estimates together with pointwise confidence intervals. We refer to Section 6 for further discussion.




Figure 6: Bivariate scatterplots of observations in the joint tail from each data set on Laplace margins. A colour gradient is used to illustrate the changing joint behaviour over time, with observations at the start and end of the time periods illustrated in dark and light blue, respectively.

To apply the framework, the same tuning parameters and formulations as suggested in Section 4.4, combined with the L^2 norm, are used for model fitting. Figure 7 illustrates the time-varying limit sets for each pair of stocks; the corresponding threshold quantile sets, i.e., $\widehat{\mathcal{R}}_t^\tau = \{\mathbf{v}(\phi)r^\tau(\phi, t) : \phi \in [0, 2\pi)\}$

are given in Appendix D. These estimates suggest a range of potentially complex, non-linear extremal dependence trends are present for each pair, with the type of trend varying over quadrants. For example, within the positive quadrant, the extremal dependence varies moderately for the Apple–Google pair but only very slightly for the remaining pairs. Our analysis also reveals the dynamics governing joint extreme losses (i.e., third quadrant). For example, the pairs Apple–Amazon and Apple–Nvidia, shown in Figure 7, exhibit strengthening extremal dependence in the lower-left (third) quadrant, suggesting that joint losses between these stocks have become more frequent in recent periods. These instances illustrate the benefits of applying the geometric approach on Laplace margins: one can evaluate the dependence structure across *all* quadrants simultaneously, providing a more complete picture of trends in extremal dependence.




Figure 7: Non-stationary limit set estimates $\widehat{\partial\mathcal{G}}_t$ over time for each of the filtered log-returns pairs'. The colour scale is used to illustrate the variation in time, with the blue and orange sets corresponding to the start and end of the observation periods, respectively.

5.4 Diagnostics

To assess the models fits, we apply the diagnostics introduced in Section 3.6. Firstly, PP and QQ plots for the truncated gamma model fits were computed with 95% confidence intervals; Figure 8 illustrates the QQ plots for all pairs, while the two variants of PP plots are given in Appendix D. One can observe generally good agreement between the observed and theoretical quantiles in all cases, although there is some deviation in the tail for the Apple–Nvidia pairing and the bodies of the Apple–Amazon and Apple–Google pairs. Overall, these plots indicate the truncated gamma model assumption of equation (3) is reasonable for capturing the radial tails across all pairs.




Figure 8: Truncated gamma QQ plot diagnostic for each pair of stocks. The red lines denote equity, and the grey shaded regions denote 95% pointwise confidence intervals.

The return level set diagnostic is also computed for each model fit. Specifically, we define a set of 200 equally spaced probabilities from 0.8 to 0.99, and evaluate the average empirical probabilities over time. Letting p and \hat{p} be the true and empirical probabilities, respectively, we compare the pairs $(-\log(1 - p), -\log(1 - \hat{p}))$; this allows us to evaluate performance for reasonably extreme return level set probabilities. The corresponding probability plots are given in Appendix D. In all cases, we observe very good agreement between the pairs, indicating the estimated return level sets (and model fits) correctly represent the structure in the data.

5.5 Simulation and evaluating joint risk

Extremal dependence across financial episodes

We now demonstrate further potential use cases for our framework. To begin, we summarise the trends in dependence via estimates of coefficients. As detailed in Appendix B, coefficients of tail dependence, which summarise the form of joint tail behaviour, can be defined separately for each quadrant. As noted in Section 2, these coefficients can be approximated directly from the estimated limit sets using the theoretical results introduced in Nolde (2014) and Nolde and Wadsworth (2022). For each pair, the time-varying η_t estimates in quadrants 1 and 3 (i.e., the joint upper and lower tails, respectively) are shown in Figure 9, with the remaining quadrants considered in Appendix D. Such coefficients provide interpretable summaries of the extremal dependence behaviour over time for different stock market dynamics. For example, these estimates suggest that the probability of joint high or low days for Apple and Nvidia returns are increasing throughout the time period, whereas the same probabilities for Apple and Google vary non-linearly over the observation period. It is also clear that the estimated trends can be very slight in some cases, such as the Apple and Microsoft pairing.

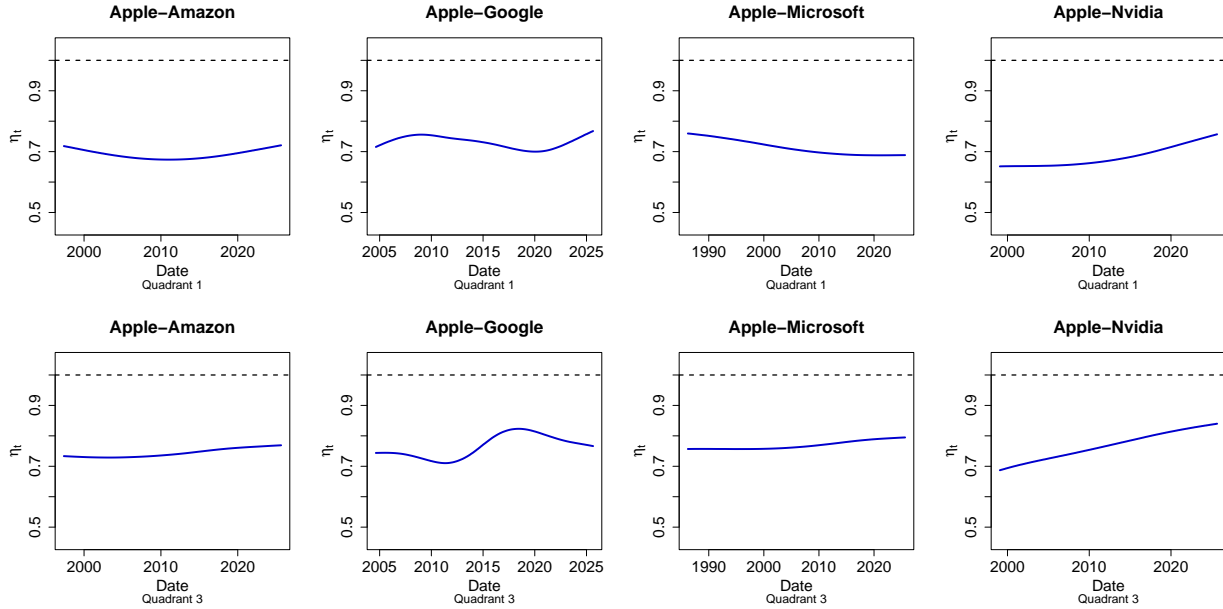


Figure 9: The coefficient of tail dependence η_t over time for quadrants 1 (top row) and 3 (bottom row). The black dotted line denotes the upper bound for η_t .

Setting $p = 0.999$, we apply the framework to obtain estimates of return level sets, denoted $\widehat{\mathcal{A}}_t^p$, which can be used as practical risk measures to inform decision making; see Haselsteiner et al. (2019), Papastathopoulos et al. (2023), and Simpson and Tawn (2024b) for further discussion. We consider three recent dates of significance in the financial context; namely, the onset of the COVID 2020 pandemic (March 16th, 2020, Mazur et al., 2021), the beginning of the Russian invasion of Ukraine (February 24th, 2022, Lo et al., 2022), and the 2025 ‘Liberation Day’ tariffs (April 2nd, 2025, Ignatenko et al., 2025). While the onset of the pandemic cannot be expressed by a single day, stock markets experienced significant volatility in mid-March, and thus this time period is of interest for studying joint tail behaviour. Return level sets for each date were computed across all pairs. The estimated sets were then transformed back to the original scale of the log-returns data, i.e., inverting the modelling procedure described in Section 5.2 and effectively ‘de-filtering’ the data. The resulting sets are illustrated in Figure 10. These plots make it clear that the behaviour of this joint risk measure varies noticeably not only between different pairings, but also over different observation dates. We note here that differences across observations dates arise from both fluctuations in the dependence structure and marginal behaviour. For example, we observe significantly more market volatility around the onset of COVID, which helps to explain the larger return level set magnitudes on this date. Since all considered dates are relatively close in proximity, we observe only subtle differences in the extremal dependence structure; this is illustrated in Appendix D, where we show the same return level set estimates on the standard Laplace scale.




Figure 10: Return level set estimates at $p = 0.999$ across all pairs for recent dates of significance in the financial context.

Generative extreme value modelling

Finally, we demonstrate how our framework can be used to simulate from the conditional variable $(\mathbf{X}_t \mid R_t > r^\tau(\phi, t))$, i.e., the random vector in regions where the radial component is relatively large. Given $t \in \{1, \dots, T\}$, this involves two steps:

1. Simulation from $(\Phi_t \mid R_t > r^\tau(\phi, t))$.
2. Simulation from $(R_t \mid \Phi_t = \phi, R_t > r^\tau(\phi, t))$.

Given a simulation ϕ from $(\Phi_t \mid R_t > r^\tau(\phi, t))$, step 2) simply involves applying inverse transform sampling using equation (3). Therefore, all that remains is to specify a model for step 1) from which angular observations can be simulated. For this, we apply a non-parametric kernel density approach known as the Nadaraya-Watson estimator (De Gooijer and Zerom, 2003). Here, we set $\hat{f}_t(\phi) = [\sum_{t^* \in \mathcal{T}^\tau} K_{h_1}^\phi(\phi - \phi_{t^*}) K_{h_2}^t(\|t^* - t\|)] / [\sum_{t^* \in \mathcal{T}^\tau} K_{h_2}^t(\|t^* - t\|)]$, where $\mathcal{T}^\tau := \{t : \|\mathbf{x}_t\| > \hat{r}^\tau(\text{atan2}(x_{2,t}, x_{1,t}), t)\}$, and $K_{h_1}^\phi$ and $K_{h_2}^t$ denote kernel functions with bandwidth parameters $h_1 > 0$ and $h_2 > 0$ respectively. We use a circular von-Mises kernel for $K_{h_1}^\phi$ and a Gaussian kernel for $K_{h_2}^t$; the former accounts for the periodicity in the angular component. Bandwidth parameters are selected manually by comparing the estimated density functions to local histograms defined on time windows. Numerical integration is then applied to simulate directly from $\hat{f}_t(\cdot)$ for any t . See Holmes et al. (2012) and Castro-Camilo et al. (2018) for similar applications of this modelling technique.

For each time point $t \in \{1, \dots, T\}$, we can simulate from the entire joint tail of $(\mathbf{X}_t \mid R_t > r^\tau(\phi, t))$. This presents a wide range of potential uses, and we present one specific use case here. When managing portfolios, practitioners often wish to assess the (tail) relationships between different stocks (McNeil et al., 2015). Joint risk measures are often used to aid with this assessment. One popular measure is known as the *conditional value-at-risk* (CoVaR) (Girardi and Ergün, 2013; Adrian and Brunnermeier, 2016). Given some




Figure 11: Downside and upside VaR and CoVaR estimates for all pairs, with the ‘stress’ events defined by the Apple stock in each case.

probability p close to 0, the *downside* CoVaR at level p is defined in the non-stationary, bivariate setting via

$$\Pr(X_{2,t} \leq \text{CoVaR}_{t,p} \mid X_{1,t} \leq \text{VaR}_{t,p}(X_{1,t})) = p. \quad (7)$$

where $\text{VaR}_{t,p}(X_{1,t}) := F_{X_{1,t}}^{-1}(p)$ is termed the *downside value-at-risk* (VaR) of $X_{1,t}$. One can view $\text{CoVaR}_{t,p}$ as the resulting stress on $X_{2,t}$ given an extreme loss for $X_{1,t}$ (Nolde and Zhou, 2021). However, as noted in Section 5.1, we are not just interested in extreme losses within our analysis. Thus, we adapt equation (7) to also consider the joint upper tail and define the *upside* CoVaR at level p via

$$\Pr(X_{2,t} \geq \text{CoVaR}_{t,p}^* \mid X_{1,t} \geq \text{VaR}_{t,p}^*(X_{1,t})) = p. \quad (8)$$

where $\text{VaR}_{t,p}^*(X_{1,t}) := F_{X_{1,t}}^{-1}(1-p)$ is termed the *upside* VaR of $X_{1,t}$. Conversely, $\text{CoVaR}_{t,p}^*$ describes the effect on $X_{2,t}$ given an extreme increase for $X_{1,t}$. We note the downside and upside naming conventions follow from similar nomenclature in the literature (e.g., Reboredo et al., 2016).

Observe that simulation can be used to approximate both $\text{CoVaR}_{t,p}$ and $\text{CoVaR}_{t,p}^*$ by empirically estimating the probabilities denoted in equations (7) and (8) and using a rootfinder to solve for equality. For each date, we simulate 100,000 observations in the joint tail and estimate the CoVaR measures at $p = 0.01$ relative to the Apple returns (i.e., $X_{1,t}$ denoting the Apple series). Figure 11 illustrates the VaR and CoVaR estimates for each pair of stocks on the observation scale; the corresponding estimates on the Laplace scale are given in Appendix D. In these plots, we also illustrate recent US business cycles relevant to the returns data in question. Each cycle corresponds to a period of recession, beginning at the peak (P) and ending at the trough (T) (see, e.g., de Carvalho et al., 2012). One can observe that the extremes in the VaR estimates for $X_{1,t}$ do not always correspond to the extremes in CoVaR, illustrating the time-varying nature of the extremal dependence within the data. This provides evidence that classical multivariate extreme value modelling approaches would not provide adequate flexibility for capturing the structure in this data.




Figure 12: Downside and upside VaR and CoVaR estimates for all pairs in the years 2018–2022, with the ‘stress’ events defined by the Apple stock in each case.

Additionally, in Appendix D, we illustrate VaR estimates of $X_{2,t}$ against CoVaR estimates on the observation scale. We observe clear correlation between these estimates, yet CoVaR always dominates, owing to the impact of the extremal dependence between the stocks.

Interestingly, we also observe in Figure 11 that both joint extreme gains and losses appear possible during recession periods. However, it is not just recession periods where such fluctuations are possible; considering the two years before and after 2020 (i.e., the COVID pandemic) we also observe non-recession periods where spikes are present in both the upper and lower tails, as illustrated in Figure 12. In some cases, certain phenomena are recognisable (e.g., the Russian invasion of Ukraine in February 2022), while other spikes cannot easily be identified with an event or recession period. This demonstrates the complex and volatile nature of evolving extremal dependence in stock markets, alongside the importance of developing flexible modelling approaches that can capture fluctuations in this dependence.

As a final remark, we note that we observe significant trends in the CoVaR estimates on both the Laplace and observed scales, even though the estimated limit sets suggest weak trends in dependence. This occurs due to the fact our simulation procedure involves models for both the angular and radial components, and we observe significant variation in the estimated angular density $\hat{f}_t(\phi)$, $\phi \in [0, 2\pi)$, over time. This emphasises the importance of modelling both components accurately in the geometric setting.

6 Discussion

In this article, we have introduced a novel geometric modelling framework for approximating non-stationary extremal dependence structures. This framework addresses certain limitations from existing works and, crucially, allows for more general and flexible modelling within the non-stationary setting. Our proposed model can capture a wide range of dependence structures, and is not limited to either AD or AI data sets. Furthermore, while we have proposed our framework in the non-stationary setting, the techniques given in

Section 3 could be adapted to obtain stationary limit set estimates, adding to a growing literature in this area. As a by-product, our framework can be used to generate random vectors in regions where the radial component is large, thereby contributing to the fast-emerging literature on generative models for multivariate extremes, which aim to simulate data across all joint tail regions; see, e.g., de Monte et al. (2025) or Wessel et al. (2025). The ability to generate, or simulate, extremal data in different tail regions is desirable beyond the financial setting; for example, in the environmental setting to assist with flood risk assessment and catastrophe modelling (Keef et al., 2013; Quinn et al., 2019).

We close the paper with some comments on open problems and future work. Firstly, as noted in Section 3.5, the boundary set estimates obtained from our approach are not required to satisfy all of the theoretical properties of boundary sets outlined in Section 1; one could in future work explore whether this rescaling could be directly incorporated into the model fitting procedure, as was the case in Murphy-Barltrop et al. (2024a) and Campbell and Wadsworth (2024).

The proposed methodology introduced in Section 3 is specifically adapted for the case where extremal dependence varies smoothly in time. However, our framework is not restricted to this setting since GAMs are highly flexible and, with appropriate formulations, can capture a wide range of covariate interactions, including discrete covariates and abrupt regime changes (e.g., changepoints). Moreover, our code is written using the `evgam` R package, meaning it can easily be adapted to capture alternative covariate interactions. Adjusting our framework to capture more general trends in extremal dependence represents an important area for future work.

Currently, the proposed approach can only be applied the bivariate setting. Flexibly modelling non-stationary extremal dependence in the general multivariate setting remains an open problem. In the general d -dimensional setting, one can still transform to polar angles; however, fitting GAMs over such angles requires more than cyclic basis splines to achieve the correct periodicity properties. Alternatively, one could try to define models directly on the $(d - 1)$ -sphere, as in Papastathopoulos et al. (2023) and Murphy-Barltrop et al. (2024a); however, this requires frameworks that can handle the collinearity that exists between the corresponding angular variables, which is not the case for many regression techniques. We also note flexible and accurate modelling of the angular distribution in higher dimensional setups remains an open problem in the stationary setting (Kakampakou and Wadsworth, 2025), although some progress has been made in recent works (e.g., de Monte et al., 2025; Wessel et al., 2025).

There is no guarantee that our proposed estimators for the non-stationary gauge functions or boundary sets will converge to the corresponding true values. While it would be desirable to have theoretical proofs of consistency, say, such results generally necessitate strict and unreasonable modelling assumptions, which themselves can be difficult to verify. We have therefore opted for a more practical treatment of our pro-

posed estimators, noting that for real data sets, one can only evaluate model performance using diagnostic procedures, such as those presented in Sections 3.6 and 5.

One notable observation from this analysis is that the choice of norm (i.e., the selected coordinate system) greatly affects the shapes of boundary set that can be represented by the modelling framework. In particular, when the boundary set exhibits ‘pointiness’ in certain regions, it may be appropriate to select a norm for which the corresponding unit ball is also pointy in these regions. However, it remains unclear how one can select this norm in practice, and future work could therefore explore whether model selection techniques could help one to perform this selection in a robust and systematic manner.

As stated in Section 5.3, best practices for establishing the existence of time-varying extremal dependence within data are yet to be developed. Constructing tests with a reasonable statistical power is implicitly more complicated in the extreme setting, since one achieves better convergence to asymptotic models as one moves further into the tail, but this vastly reduces the available sample size and hence power. Developing robust and reliable techniques for diagnosing trends in extremal dependence represents an important line for future work.

Finally, an obvious omission from the case study of Section 5 is the lack of uncertainty quantification for the fitted models and computed statistics. Standard bootstrapping techniques are not applicable here, since resampling such data samples would not preserve a time-varying dependence structure. Uncertainty quantification is further complicated by the fact our modelling framework contains many components, each of which is fitted independently; the quantile function, the gauge function, and the angular model. Jointly quantifying the uncertainty arising from each of these components is non-trivial and would require careful consideration. Being able to address the significance of trends represents an important line for future research. Noting that the trends we found in Section 5 were very slight, we also fit a stationary version of our model. The diagnostics for this also appear satisfactory, suggesting that the observed trends may indeed not be statistically significant in this case. A starting point for considering uncertainty quantification could be to adapt established approaches for non-stationary, univariate time series, such as the local block bootstrap (e.g., Politis et al., 1999, Chapter 4).

Declarations

Ethical Approval

Not Applicable

Availability of supporting data

The data sets analysed in Section 5 are available at <https://github.com/callumbarltrop/NSGE>.

Competing interests

The authors have no relevant financial or non-financial interests to disclose.

Funding details

CMB acknowledges the financial support by the Federal Ministry of Research, Technology and Space of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI. JW gratefully acknowledges support from UK Engineering and Physical Sciences Research Council grant EP/X010449/1. MdC is partially funded by Leverhulme Trust and FCT under Grants UIDB/04106/2020 and UIDP/04106/2020.

Code availability

Code for fitting the proposed framework is available at <https://github.com/callumbarltrop/NSGE>.

Authors' contribution

CMB proposed the methodological framework, conducted the simulation study and case analyses, prepared the initial manuscript draft, and led the overall project. JW and MdC provided substantial guidance throughout the development of the work, contributed to the refinement of the methodology, and gave extensive feedback on successive drafts. BY contributed to the implementation and development of the computational code and also provided feedback on the manuscript. All authors approved the final version of the manuscript.

References

- Adrian, T. and Brunnermeier, M. K. (2016). CoVaR. *The American Economic Review*, 106(7):1705.
- André, L. M., Campbell, R., D'Arcy, E., Farrell, A., Healy, D., Kakampakou, L., Murphy, C., Murphy-Barltrop, C. J. R., and Speers, M. (2025). Extreme value methods for estimating rare events in utopia: Eva (2023) conference data challenge: team lancopula utopiversity. *Extremes*, 28(1):23–45.
- Balkema, G. and Nolde, N. (2010). Asymptotic independence for unimodal densities. *Advances in Applied Probability*, 42:411–432.

- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327.
- Brunner, M. I., Seibert, J., and Favre, A. C. (2016). Bivariate return periods and their importance for flood peak and volume estimation. *Wiley Interdisciplinary Reviews: Water*, 3:819–833.
- Campbell, R. and Wadsworth, J. (2024). Piecewise-linear modeling of multivariate geometric extremes. *arXiv*, 2412.05195.
- Castro-Camilo, D., de Carvalho, M., and Wadsworth, J. (2018). Time-varying extreme value dependence with application to leading European stock markets. *Annals of Applied Statistics*, 12:283–309.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer, London.
- Coles, S., Heffernan, J., and Tawn, J. (1999). Dependence measures for multivariate extremes. *Extremes*, 2:339–365.
- Coles, S. G. and Tawn, J. A. (1991). Modelling extreme multivariate events. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 53:377–392.
- David, H. A. and Nagaraja, H. N. (2004). *Order Statistics*. Wiley, New York.
- Davis, R. A. and Mikosch, T. (2009). The extremogram: A correlogram for extreme events. *Bernoulli*, 15(4):977–1009.
- Davis, R. A., Mulrow, E., and Resnick, S. I. (1988). Almost sure limit sets of random samples in \mathbb{R}^d . *Advances in Applied Probability*, 20:573–599.
- de Carvalho, M. and Davison, A. C. (2014). Spectral density ratio models for multivariate extremes. *Journal of the American Statistical Association*, 109:764–776.
- de Carvalho, M., Huser, R., Naveau, P., and Reich, B. J. (2026). *Handbook of Statistics of Extremes*. Chapman & Hall/CRC, Boca Raton, FL.
- de Carvalho, M. and Palacios Ramirez, K. V. (2025). Semiparametric bayesian modelling of nonstationary joint extremes: How do big tech’s extreme losses behave? *Journal of the Royal Statistical Society, Series C: Applied Statistics*, 74(2):447–465.
- de Carvalho, M., Rodrigues, P. C., and Rua, A. (2012). Tracking the US business cycle with a singular spectrum analysis. *Economics letters*, 114(1):32–35.
- De Gooijer, J. G. and Zerom, D. (2003). On conditional density estimation. *Statistica Neerlandica*, 57(2):159–176.
- de Monte, L., Huser, R., Papastathopoulos, I., and Richards, J. (2025). Generative modelling of multivariate geometric extremes using normalising flows. *arXiv*, 2505.02957.
- Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74:427–431.

- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50:987.
- Fasiolo, M., Wood, S. N., Zaffran, M., Nedellec, R., and Goude, Y. (2021a). Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 116:1402–1412.
- Fasiolo, M., Wood, S. N., Zaffran, M., Nedellec, R., and Goude, Y. (2021b). qgam : Bayesian nonparametric quantile regression modeling in R. *Journal of Statistical Software*, 100.
- Frolova, N. and Cribben, I. (2016). *Extremogram: Estimation of Extreme Value Dependence for Time Series Data*. Comprehensive R Archive Network (CRAN). R package version 1.0.2.
- Girardi, G. and Ergün, A. T. (2013). Systemic risk measurement: Multivariate garch estimation of covar. *Journal of Banking & Finance*, 37(8):3169–3180.
- Gong, Y. and Huser, R. (2022). Asymmetric tail dependence modeling, with application to cryptocurrency market data. *The Annals of Applied Statistics*, 16:1822–1847.
- Gouldby, B., Wyncoll, D., Panzeri, M., Franklin, M., Hunt, T., Hames, D., Tozer, N., Hawkes, P., Dornbusch, U., and Pullen, T. (2017). Multivariate extreme value modelling of sea conditions around the coast of England. *Proceedings of the Institution of Civil Engineers: Maritime Engineering*, 170:3–20.
- Guerrero, M. B., Huser, R., and Ombao, H. (2023). Conex–Connect: Learning patterns in extremal brain connectivity from MultiChannel EEG data. *The Annals of Applied Statistics*, 17:178–198.
- Haselsteiner, A. F., Nguyen, P. T., Coe, R. G., Martin, N., Manuel, L., and Eckert-Gallup, A. (2019). A benchmarking exercise on estimating extreme environmental conditions: Methodology and baseline results. *Proceedings of the International Conference on Offshore Mechanics and Arctic Engineering - OMAE*, 3:1–10.
- Haver, S. and Kleiven, G. (2004). Environmental contour lines for design purposes - why and when? *Proceedings of the International Conference on Offshore Mechanics and Arctic Engineering - OMAE*, 1:337–345.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 66:497–546.
- Holmes, M. P., Gray, A. G., and Isbell, C. L. (2012). Fast nonparametric conditional density estimation. *arXiv*, 1206.5278.
- Huser, R., Opitz, T., and Wadsworth, J. L. (2025). Modeling of spatial extremes in environmental data science: Time to move away from max-stable processes. *Environmental Data Science*, 4:e3.
- Huser, R. and Wadsworth, J. L. (2019). Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 114:434–444.

- Ignatenko, A., Lashkaripour, A., Macedoni, L., and Simonovska, I. (2025). Making America great again? The economic impacts of liberation day tariffs. Technical report, National Bureau of Economic Research.
- Jonathan, P., Ewans, K., and Flynn, J. (2014a). On the estimation of ocean engineering design contours. *Journal of Offshore Mechanics and Arctic Engineering*, 136:1–8.
- Jonathan, P., Ewans, K., and Randell, D. (2014b). Non-stationary conditional extremes of northern North Sea storm characteristics. *Environmetrics*, 25:172–188.
- Kakampakou, L., Simpson, E. S., and Wadsworth, J. L. (2024). Spatial extremal modelling: A case study on the interplay between margins and dependence. *Stat*, 13(4):e70021.
- Kakampakou, L. and Wadsworth, J. L. (2025). Geometric modelling of spatial extremes. *arXiv preprint*, 2511.08192.
- Keef, C., Tawn, J. A., and Lamb, R. (2013). Estimating the probability of widespread flood events. *Environmetrics*, 24:13–21.
- Kinoshita, K. and Resnick, S. I. (1991). Convergence of scaled random samples in \mathbb{R}^d . *The Annals of Probability*, 19:1640–1663.
- Koenker, R. (2011). Additive models for quantile regression: Model selection and confidence band-aids. *Brazilian Journal of Probability and Statistics*, 25:239–262.
- Koenker, R., Chernozhukov, V., He, X., and Peng, L. (2017). *Handbook of Quantile Regression*. Chapman & Hall/CRC, Boca Raton, FL.
- Kwiatkowski, D., Phillips, P. C., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54:159–178.
- Ledford, A. W. and Tawn, J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83:169–187.
- Ledford, A. W. and Tawn, J. A. (1997). Modelling dependence within joint tail regions. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 59:475–499.
- Lee, J., de Carvalho, M., Rua, A., and Avila, J. (2024). Bayesian smoothing for time-varying extremal dependence. *Journal of the Royal Statistical Society, Series C: Applied Statistics*, 73:581–597.
- Liu, Y. and Tawn, J. A. (2014). Self-consistent estimation of conditional multivariate extreme value distributions. *Journal of Multivariate Analysis*, 127:19–35.
- Lo, G.-D., Marcelin, I., Bassène, T., and Sène, B. (2022). The Russo-Ukrainian war and financial markets: the role of dependence on Russian commodities. *Finance Research Letters*, 50:103194.

- Mackay, E. and Haselsteiner, A. F. (2021). Marginal and total exceedance probabilities of environmental contours. *Marine Structures*, 75:1–24.
- Mackay, E. and Jonathan, P. (2023). Modelling multivariate extremes through angular-radial decomposition of the density function. *arXiv*, 2310.12711.
- Mackay, E. B., Murphy-Barltrop, C. J., Richards, J., and Jonathan, P. (2025). Deep learning joint extremes of metocean variables using the SPAR model. In *International Conference on Offshore Mechanics and Arctic Engineering*, volume 88926, page V003T06A034. American Society of Mechanical Engineers.
- Majumder, R., Shaby, B. A., Reich, B. J., and Cooley, D. S. (2025). Semiparametric estimation of the shape of the limiting bivariate point cloud. *Bayesian Analysis*, 1(1):1–27.
- Mazur, M., Dang, M., and Vega, M. (2021). COVID-19 and the March 2020 stock market crash. evidence from S&P1500. *Finance Research Letters*, 38:101690.
- McNeil, A. J., Frey, R., and Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools - Revised Edition*. Princeton University Press.
- Mhalla, L., Chavez-Demoulin, V., and Naveau, P. (2017). Non-linear models for extremal dependence. *Journal of Multivariate Analysis*, 159:49–66.
- Mhalla, L., Opitz, T., and Chavez-Demoulin, V. (2019). Exceedance-based nonlinear regression of tail dependence. *Extremes*, 22:523–552.
- Murphy, C., Tawn, J. A., and Varty, Z. (2025). Automated threshold selection and associated inference uncertainty for univariate extremes. *Technometrics*, 67(2):215–224.
- Murphy-Barltrop, C. and Wadsworth, J. (2024). Modelling non-stationarity in asymptotically independent extremes. *Computational Statistics and Data Analysis*, 199:108025.
- Murphy-Barltrop, C. J. R., Mackay, E., and Jonathan, P. (2025). Inference for bivariate extremes via a semi-parametric angular-radial model. *Extremes*, 28(2):209–238.
- Murphy-Barltrop, C. J. R., Majumder, R., and Richards, J. (2024a). Deep learning of multivariate extremes via a geometric representation. *arXiv*, 2406.19936.
- Murphy-Barltrop, C. J. R., Wadsworth, J. L., and Eastoe, E. F. (2024b). Improving estimation for asymptotically independent bivariate extremes via global estimators for the angular dependence function. *Extremes*, 27:643–671.
- Nolde, N. (2014). Geometric interpretation of the residual dependence coefficient. *Journal of Multivariate Analysis*, 123:85–95.

- Nolde, N. and Wadsworth, J. L. (2022). Linking representations for multivariate extremes via a limit set. *Advances in Applied Probability*, 54:688–717.
- Nolde, N. and Zhou, C. (2021). Extreme value analysis for financial risk management. *Annual Review of Statistics and Its Application*, 8:217–240.
- Oh, H.-S., Lee, T. C. M., and Nychka, D. W. (2011). Fast nonparametric quantile regression with arbitrary smoothing methods. *Journal of Computational and Graphical Statistics*, 20:510–526.
- Opitz, T. (2016). Modeling asymptotically independent spatial extremes based on Laplace random fields. *Spatial Statistics*, 16:1–18.
- Papastathopoulos, I., de Monte, L., Campbell, R., and Rue, H. (2023). Statistical inference for radially-stable generalized pareto distributions and return level-sets in geometric extremes. *arXiv*, 2310.06130.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2):527–556.
- Peprah, W. K., Rafaliniony, H., and Peprah, P. (2025). A longitudinal study of the impact of the covid-19 pandemic on the managerial effectiveness of the “magnificent seven” stocks firms: An economic value-added approach. *Review of Pacific Basin Financial Markets and Policies*, 28(01):2450035.
- Politis, D. N., Romano, J. P., and Wolf, M. (1999). *Subsampling*. Springer, New York.
- Poon, S.-H., Rockinger, M., and Tawn, J. (2003). Modelling extreme-value dependence in international stock markets. *Statistica Sinica*, pages 929–953.
- Quinn, N., Bates, P. D., Neal, J., Smith, A., Wing, O., Sampson, C., Smith, J., and Heffernan, J. (2019). The spatial dependence of flood hazard and risk in the United States. *Water Resources Research*, 55:1890–1911.
- Reboredo, J. C., Rivera-Castro, M. A., and Ugolini, A. (2016). Downside and upside risk spillovers between exchange rates and stock prices. *Journal of Banking & Finance*, 62:76–96.
- Salvatierra, I. D. L. and Patton, A. J. (2015). Dynamic copula models and high frequency data. *Journal of Empirical Finance*, 30:120–135.
- Simpson, E. S. and Tawn, J. A. (2024a). Estimating the limiting shape of bivariate scaled sample clouds: With additional benefits of self-consistent inference for existing extremal dependence properties. *Electronic Journal of Statistics*, 18(2):4582–4611.
- Simpson, E. S. and Tawn, J. A. (2024b). Inference for new environmental contours using extreme value analysis. *Journal of Agricultural, Biological and Environmental Statistics*, 30:638–662.

- Simpson, E. S., Wadsworth, J. L., and Tawn, J. A. (2020). Determining the dependence structure of multivariate extremes. *Biometrika*, 107:513–532.
- Sklar, A. (1959). Fonctions de repartition a n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231.
- Talento, M. S. D., Richards, J., Pinto-Orellana, M., Huser, R., and Ombao, H. C. (2025). Spectral extremal connectivity of two-state seizure brain waves. *arXiv*, 2503.04169.
- Tendijck, S., Eastoe, E., Tawn, J., Randell, D., and Jonathan, P. (2021). Modeling the extremes of bivariate mixture distributions with application to oceanographic data. *Journal of the American Statistical Association*, 118:1373–1384.
- Tsay, R. S. (2010). *Analysis of Financial Time Series*. Wiley, New York.
- Wadsworth, J. L. and Campbell, R. (2024). Statistical inference for multivariate extremes via a geometric approach. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 86:1243–1265.
- Wadsworth, J. L. and Tawn, J. A. (2013). A new representation for multivariate tail probabilities. *Bernoulli*, 19:2689–2714.
- Wadsworth, J. L., Tawn, J. A., Davison, A. C., and Elton, D. M. (2017). Modelling across extremal dependence classes. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 79:149–175.
- Wessel, J. B., Murphy-Barltrop, C. J., and Simpson, E. S. (2025). Generative machine learning for multivariate angular simulation. *Extremes*, pages 1–49.
- Wood, S. N. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 65:95–114.
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 73:3–36.
- Wood, S. N. (2017). *Generalized Additive Models*. Chapman & Hall/CRC, Boca Raton, FL.
- Wood, S. N., Pya, N., and Säfken, B. (2016). Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association*, 111:1548–1563.
- Yoshida, T. (2026). Asymptotic theory for extreme value generalized additive model. *Journal of Statistical Theory and Practice*, 20(1):6.
- Youngman, B. D. (2019). Generalized additive models for exceedances of high thresholds with an application to return level estimation for U.S. wind gusts. *Journal of the American Statistical Association*, 114:1865–1879.
- Youngman, B. D. (2022). evgam: An R package for generalized additive extreme value models. *Journal of Statistical Software*, 103.

Yu, K. and Moyeed, R. A. (2001). Bayesian quantile regression. *Statistics and Probability Letters*, 54:437–447.

Appendix

A Additional illustrative figures

Figure A illustrates the function $\rho_2(t)$ defined in Section 4.1.




Figure A: Plot illustrating the correlation parameter function $\rho_2(t)$ over t for the second copula example of Section 4.1.

Figure B illustrates two choices of knot locations for the polar angle Φ_t . Observe that both sets of knots intersect all axes and primary diagonals.

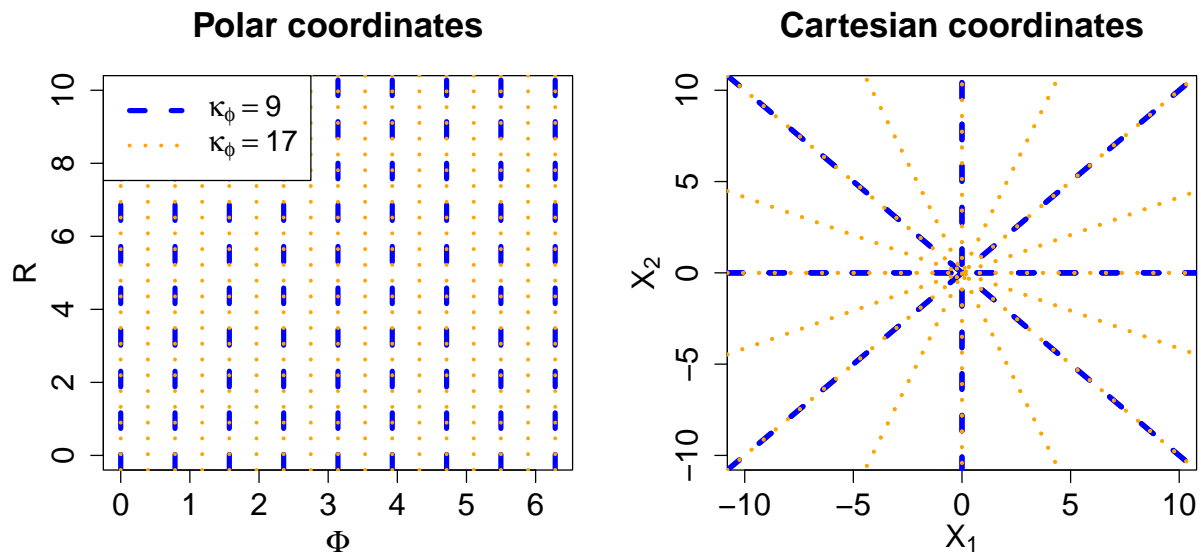


Figure B: Figure illustrating the two choices of knot locations for Φ_t in polar (left) and Cartesian (right) coordinates.

Figure illustrates tensor product basis splines at two pairs of knots for the predictor variables ϕ and t , with $T = 100$. One can observe the cyclic nature of the splines at the end points of the angular component.




Figure C: Example tensor product basis functions for ϕ and t , with $T = 100$.

B The Ledford and Tawn (1996) model on standard Laplace margins.

While originally proposed on unit Fréchet margins, the Ledford and Tawn (1996) model can be generalised to other marginal scales (see, e.g., Simpson and Tawn, 2024a; Murphy-Barltrop et al., 2024a). Given a random vector $\mathbf{X} = (X_1, X_2)$ on standard Laplace margins, we consider the following modelling formulation: for any $\boldsymbol{o} = (o_1, o_2) \in \{-1, 1\}^2$, we assume

$$\Pr(\boldsymbol{o}\mathbf{X} > u) \sim L(e^u) \exp(-u/\eta_{\boldsymbol{o}}) \text{ as } u \rightarrow \infty, \quad (\text{A.1})$$

where $L(\cdot)$ is a slowly varying function, i.e., $\lim_{u \rightarrow \infty} L(cu)/L(u) = 1$ for any constant $c > 0$, and $\eta_{\boldsymbol{o}} \in (0, 1]$ is termed the coefficient of tail dependence. These coefficients provide information about the tail dependence structure; for example, setting $\boldsymbol{o} = (1, 1)$, asymptotic dependence corresponds to the case when $\eta = 1$ and $\lim_{u \rightarrow \infty} L(u) > 0$. Observe also that $\boldsymbol{o} = (1, 1)$, $\boldsymbol{o} = (-1, 1)$, $\boldsymbol{o} = (-1, -1)$, and $\boldsymbol{o} = (1, -1)$ correspond

to the first, second, third, and fourth quadrants, or orthants, respectively. For simplicity, we drop the $(\cdot)_o$ notation throughout the article, making it clear in each instance which quadrant we are referring to.

C Simulation study results

In this section, we provide figures illustrating the results of the simulation study detailed in Section 4. We present these figures in several subsections, corresponding to which component of the model formulation we are considering. Note that the legend for each figure is identical to that of Figure 4 in the main text.

C.1 Evaluating the effect of quantile level

Figures D-AL illustrate the effect of the quantile level τ on the boundary set estimates.




Figure D: Boundary set estimates as $t = 1$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.




Figure E: Boundary set estimates as $t = T/2$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.

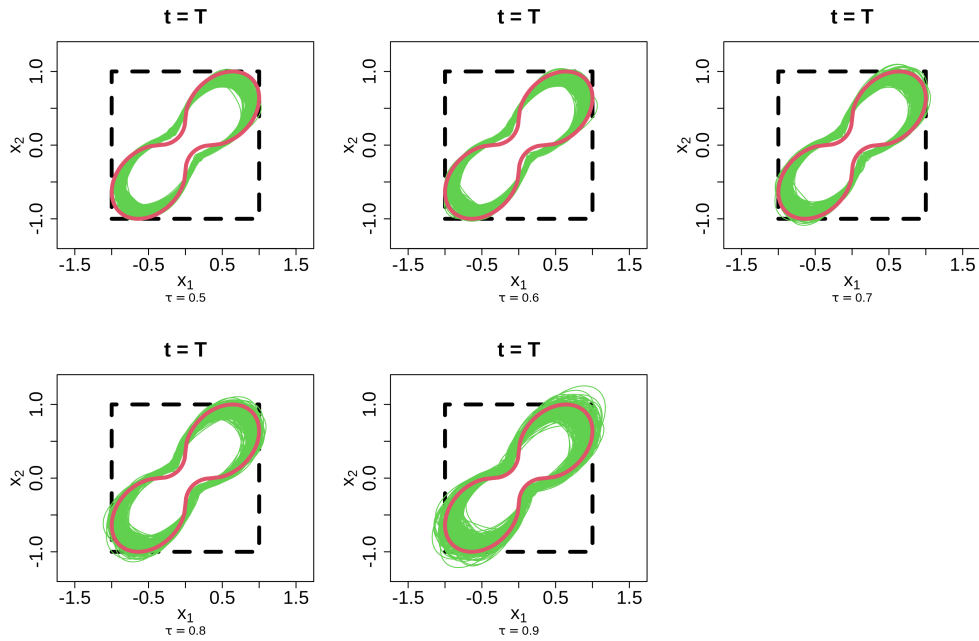


Figure F: Boundary set estimates at $t = T$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.

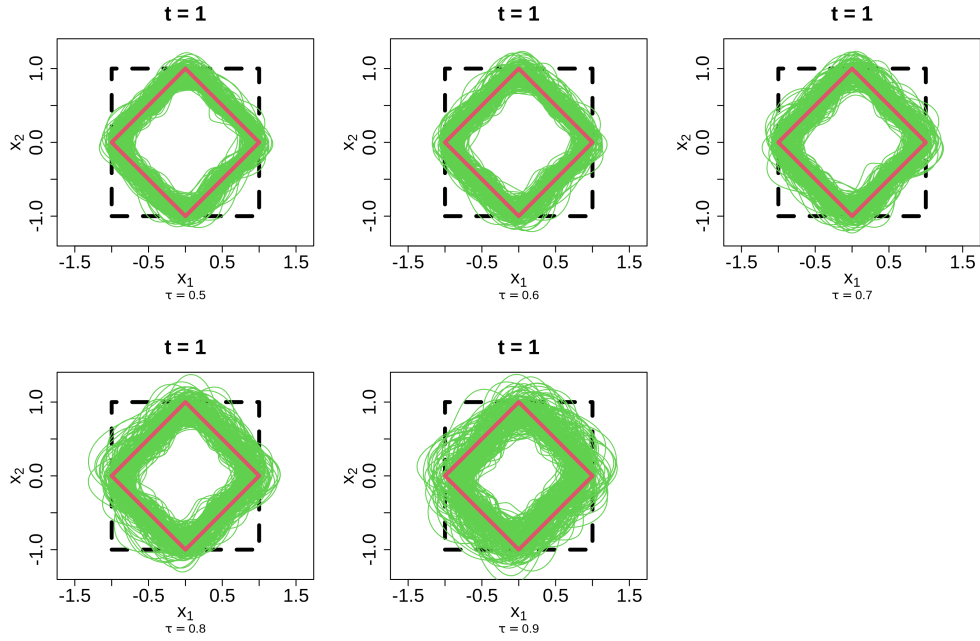


Figure G: Boundary set estimates as $t = 1$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.




Figure H: Boundary set estimates as $t = T/2$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.




Figure I: Boundary set estimates at $t = T$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.

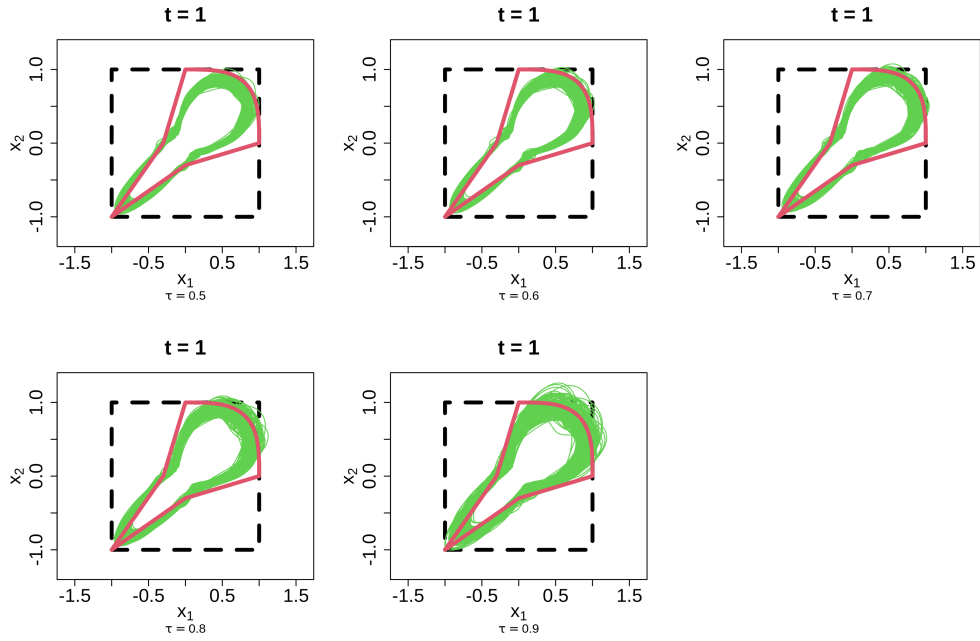


Figure J: Boundary set estimates as $t = 1$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.

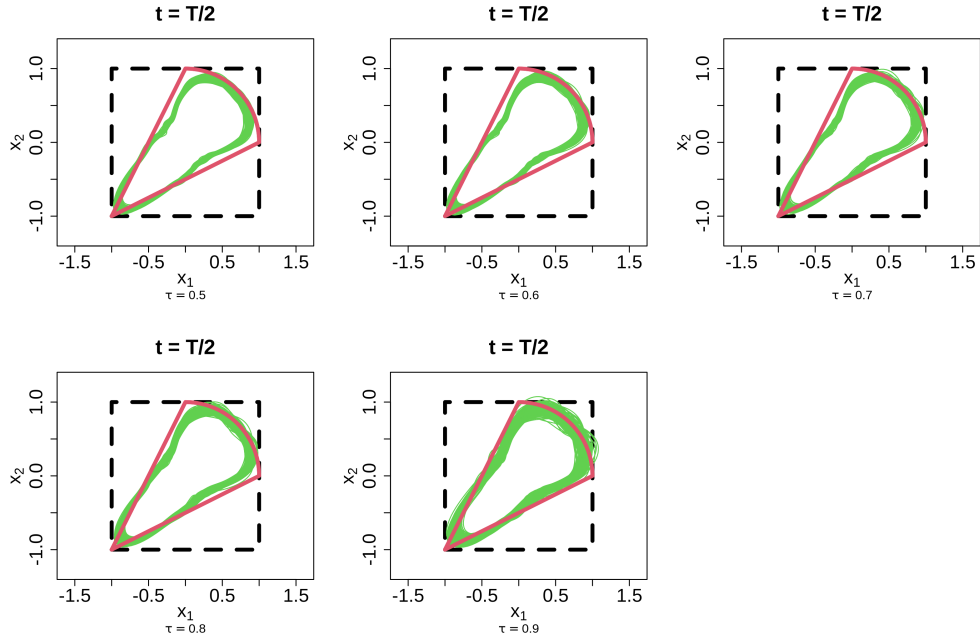


Figure K: Boundary set estimates as $t = T/2$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.

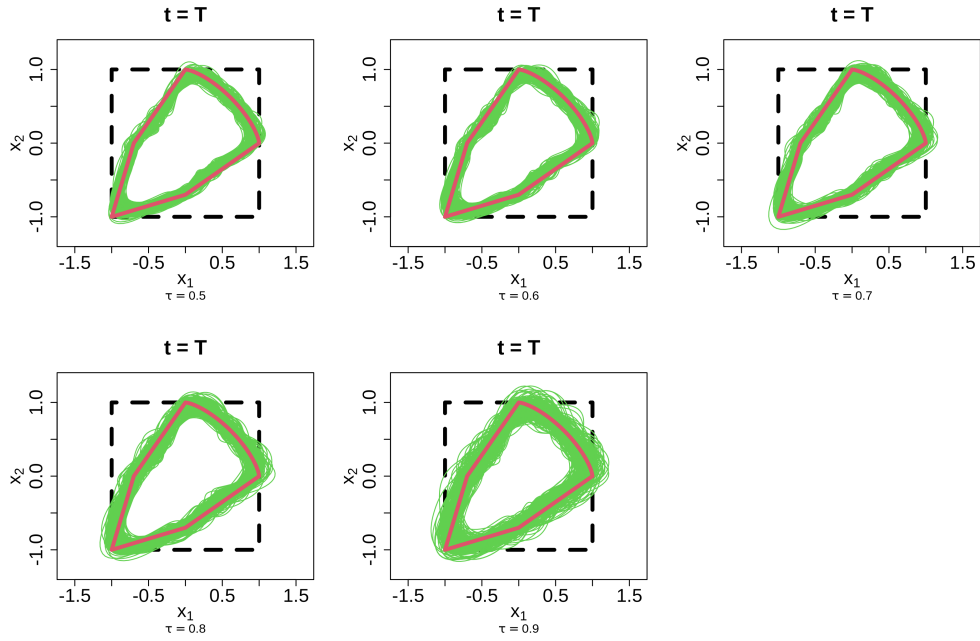


Figure L: Boundary set estimates at $t = T$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.

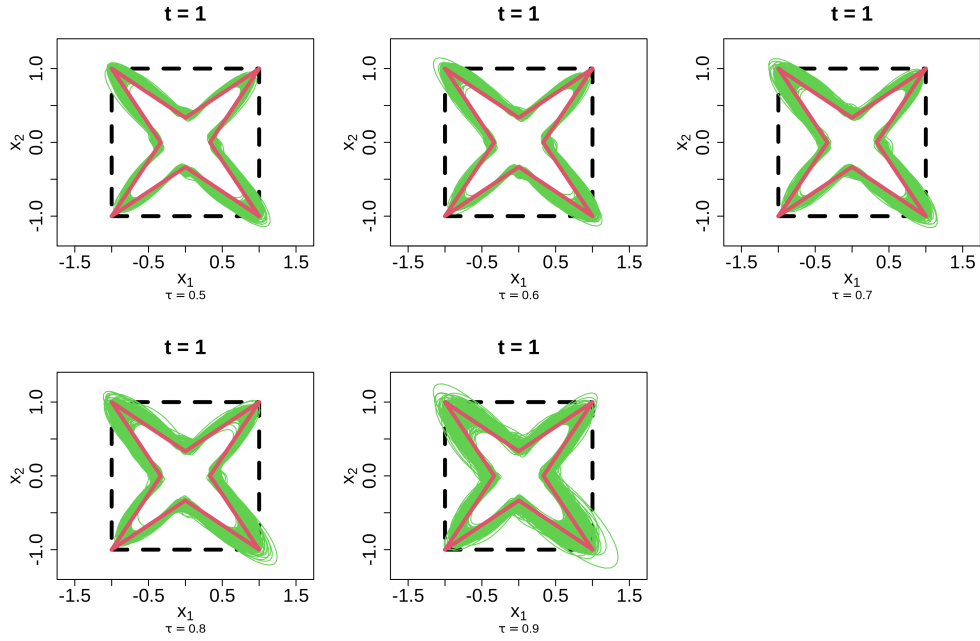


Figure M: Boundary set estimates as $t = 1$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.




Figure N: Boundary set estimates as $t = T/2$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.




Figure O: Boundary set estimates at $t = T$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.

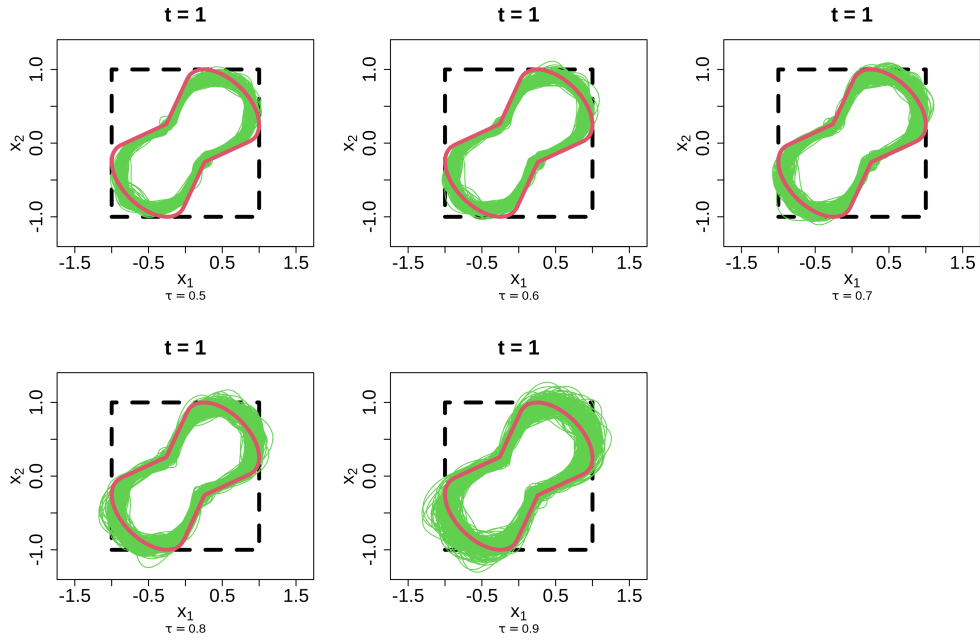


Figure P: Boundary set estimates as $t = 1$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.

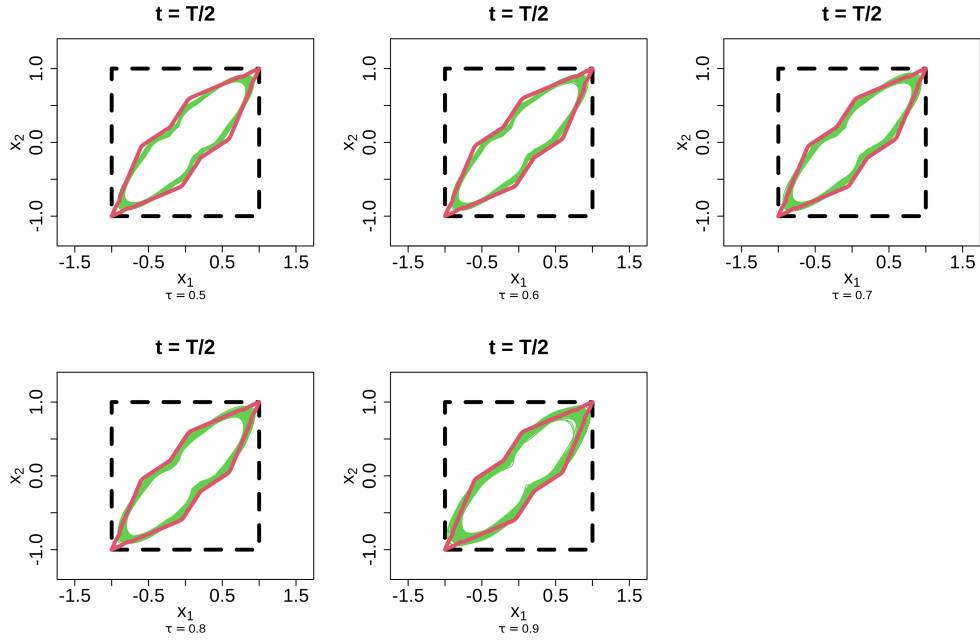


Figure Q: Boundary set estimates as $t = T/2$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.

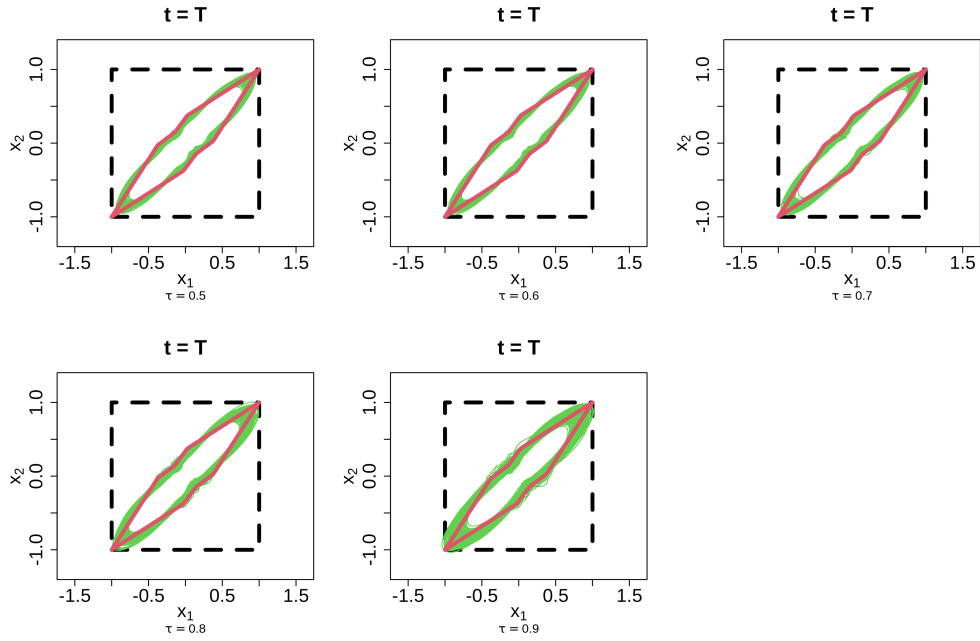


Figure R: Boundary set estimates at $t = T$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.




Figure S: Boundary set radii estimates at $\phi = \pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.




Figure T: Boundary set radii estimates at $\phi = 3\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.




Figure U: Boundary set radii estimates at $\phi = 5\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.




Figure V: Boundary set radii estimates at $\phi = 7\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the first copula example.




Figure W: Boundary set radii estimates at $\phi = \pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.




Figure X: Boundary set radii estimates at $\phi = 3\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.




Figure Y: Boundary set radii estimates at $\phi = 5\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.




Figure Z: Boundary set radii estimates at $\phi = 7\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the second copula example.




Figure AA: Boundary set radii estimates at $\phi = \pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.




Figure AB: Boundary set radii estimates at $\phi = 3\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.




Figure AC: Boundary set radii estimates at $\phi = 5\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.




Figure AD: Boundary set radii estimates at $\phi = 7\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the third copula example.




Figure AE: Boundary set radii estimates at $\phi = \pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.




Figure AF: Boundary set radii estimates at $\phi = 3\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.




Figure AG: Boundary set radii estimates at $\phi = 5\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.




Figure AH: Boundary set radii estimates at $\phi = 7\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fourth copula example.




Figure AI: Boundary set radii estimates at $\phi = \pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.




Figure AJ: Boundary set radii estimates at $\phi = 3\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.




Figure AK: Boundary set radii estimates at $\phi = 5\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.




Figure AL: Boundary set radii estimates at $\phi = 7\pi/4$ across $\tau \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the fifth copula example.

C.2 Evaluating the effect of basis dimension

Figures AM-BU illustrate the effect of the basis dimension κ_t on the boundary set estimates, while Figures BV-DD illustrate the effect of the basis dimension κ_ϕ .




Figure AM: Boundary set estimates as $t = 1$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure AN: Boundary set estimates as $t = T/2$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure AO: Boundary set estimates at $t = T$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure AP: Boundary set estimates as $t = 1$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure AQ: Boundary set estimates as $t = T/2$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure AR: Boundary set estimates at $t = T$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure AS: Boundary set estimates as $t = 1$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure AT: Boundary set estimates as $t = T/2$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure AU: Boundary set estimates at $t = T$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure AV: Boundary set estimates as $t = 1$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure AW: Boundary set estimates as $t = T/2$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure AX: Boundary set estimates at $t = T$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure AY: Boundary set estimates as $t = 1$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure AZ: Boundary set estimates as $t = T/2$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure BA: Boundary set estimates at $t = T$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure BB: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure BC: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure BD: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure BE: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the first copula example.




Figure BF: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure BG: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure BH: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure BI: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the second copula example.




Figure BJ: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure BK: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure BL: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure BM: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the third copula example.




Figure BN: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure BO: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure BP: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure BQ: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fourth copula example.




Figure BR: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure BS: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure BT: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure BU: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_t \in \{5, 10, 15\}$ for the fifth copula example.




Figure BV: Boundary set estimates as $t = 1$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure BW: Boundary set estimates as $t = T/2$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure BX: Boundary set estimates at $t = T$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure BY: Boundary set estimates as $t = 1$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure BZ: Boundary set estimates as $t = T/2$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure CA: Boundary set estimates at $t = T$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure CB: Boundary set estimates as $t = 1$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CC: Boundary set estimates as $t = T/2$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CD: Boundary set estimates at $t = T$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CE: Boundary set estimates as $t = 1$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure CF: Boundary set estimates as $t = T/2$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure CG: Boundary set estimates at $t = T$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure CH: Boundary set estimates as $t = 1$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.




Figure CI: Boundary set estimates as $t = T/2$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.




Figure CJ: Boundary set estimates at $t = T$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.




Figure CK: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure CL: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure CM: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure CN: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the first copula example.




Figure CO: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure CP: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure CQ: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure CR: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the second copula example.




Figure CS: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CT: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CU: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CV: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the third copula example.




Figure CW: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure CX: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure CY: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure CZ: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fourth copula example.




Figure DA: Boundary set radii estimates at $\phi = \pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.




Figure DB: Boundary set radii estimates at $\phi = 3\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.




Figure DC: Boundary set radii estimates at $\phi = 5\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.




Figure DD: Boundary set radii estimates at $\phi = 7\pi/4$ across $\kappa_\phi \in \{9, 17\}$ for the fifth copula example.

C.3 Evaluating the effect of the norm choice

Figures DE-EM illustrate the effect of the choice of norm $\|\cdot\|$ on the boundary set estimates.




Figure DE: Boundary set estimates as $t = 1$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DF: Boundary set estimates as $t = T/2$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DG: Boundary set estimates at $t = T$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DH: Boundary set estimates as $t = 1$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure DI: Boundary set estimates as $t = T/2$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure DJ: Boundary set estimates at $t = T$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure DK: Boundary set estimates as $t = 1$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure DL: Boundary set estimates as $t = T/2$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure DM: Boundary set estimates at $t = T$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure DN: Boundary set estimates as $t = 1$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure DO: Boundary set estimates as $t = T/2$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure DP: Boundary set estimates at $t = T$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure DQ: Boundary set estimates as $t = 1$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.




Figure DR: Boundary set estimates as $t = T/2$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.




Figure DS: Boundary set estimates at $t = T$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.




Figure DT: Boundary set radii estimates at $\phi = \pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DU: Boundary set radii estimates at $\phi = 3\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DV: Boundary set radii estimates at $\phi = 5\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DW: Boundary set radii estimates at $\phi = 7\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the first copula example.




Figure DX: Boundary set radii estimates at $\phi = \pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure DY: Boundary set radii estimates at $\phi = 3\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure DZ: Boundary set radii estimates at $\phi = 5\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure EA: Boundary set radii estimates at $\phi = 7\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the second copula example.




Figure EB: Boundary set radii estimates at $\phi = \pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure EC: Boundary set radii estimates at $\phi = 3\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure ED: Boundary set radii estimates at $\phi = 5\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure EE: Boundary set radii estimates at $\phi = 7\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the third copula example.




Figure EF: Boundary set radii estimates at $\phi = \pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure EG: Boundary set radii estimates at $\phi = 3\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure EH: Boundary set radii estimates at $\phi = 5\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure EI: Boundary set radii estimates at $\phi = 7\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fourth copula example.




Figure EJ: Boundary set radii estimates at $\phi = \pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.




Figure EK: Boundary set radii estimates at $\phi = 3\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.




Figure EL: Boundary set radii estimates at $\phi = 5\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.




Figure EM: Boundary set radii estimates at $\phi = 7\pi/4$ across $\|\cdot\| \in \{\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty\}$ for the fifth copula example.

C.4 Additional results

Figures ER-FF illustrate additional results for the model formulation selected in Section 4.4.




Figure EN: Boundary set estimates for $t = T/2$. The left, centre and right columns correspond to the second, fourth and fifth copula examples, respectively.

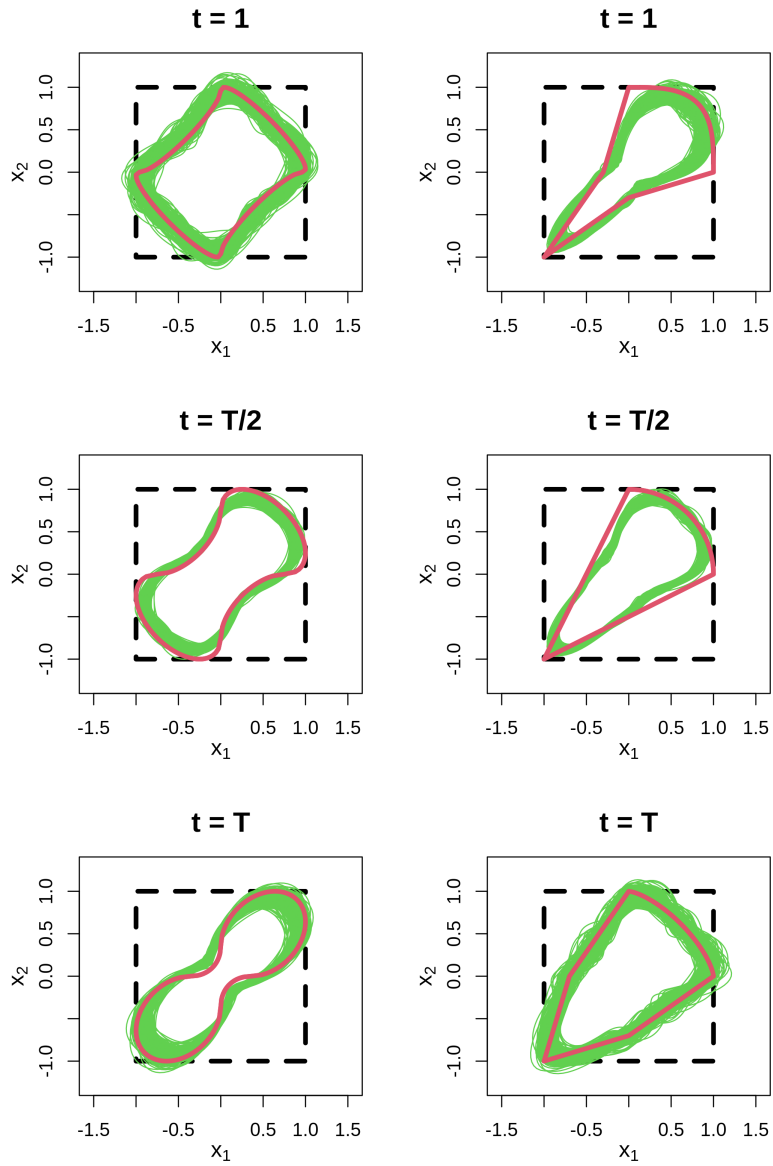


Figure EO: Boundary set estimates for $t = 1$ (top row), $t = T/2$ (middle row) and $t = T$ (bottom row). The left, centre and right columns correspond to the first and third copula examples, respectively.

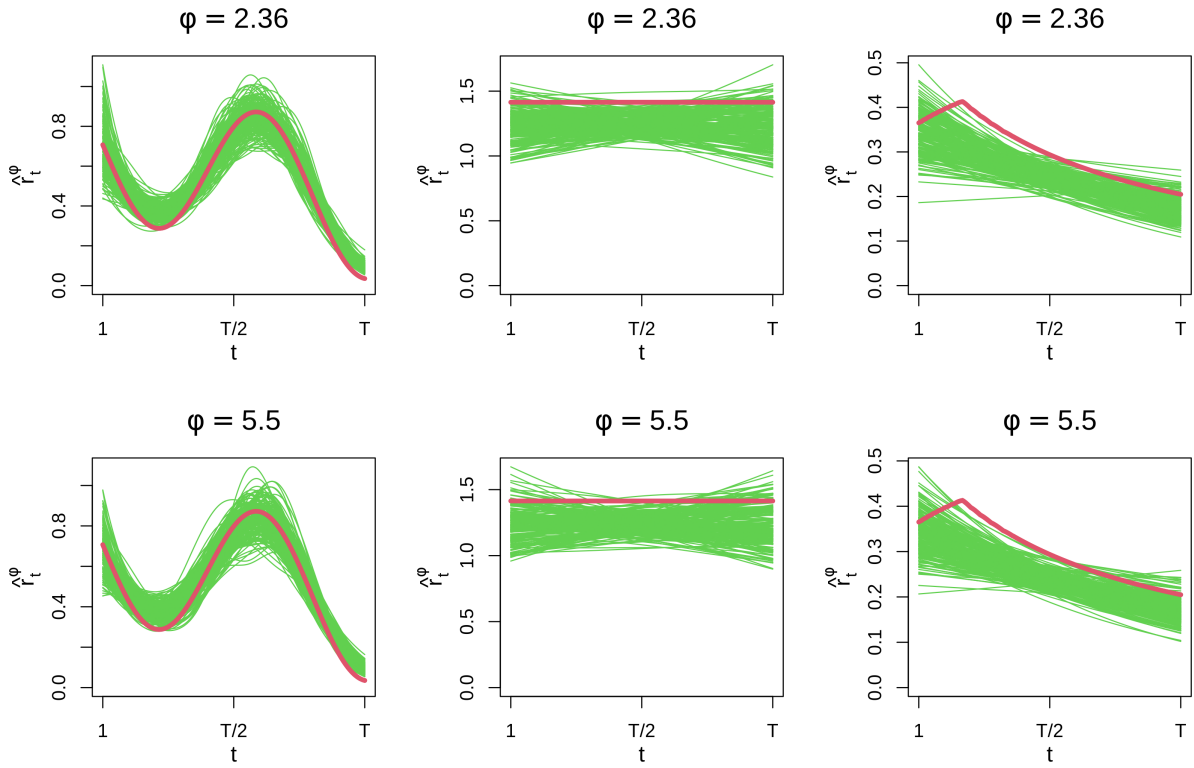


Figure EP: Boundary set radii estimates over time for $\phi = 5\pi/4$ (top row) and $\phi = 7\pi/4$ (bottom row). The left, centre and right columns correspond to the second, fourth and fifth copula examples, respectively.

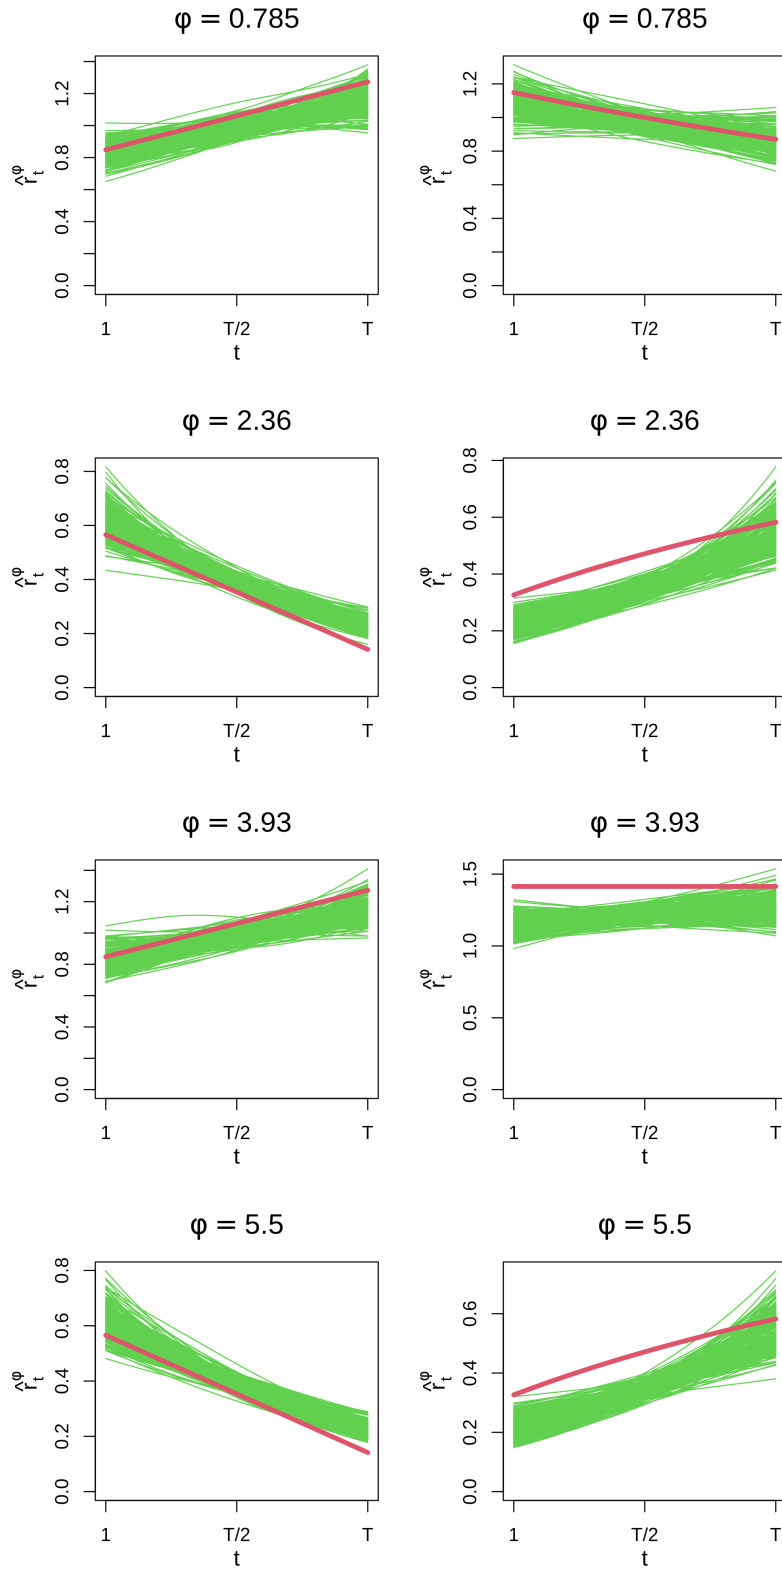


Figure EQ: Boundary set radii estimates over time for $\phi = \pi/4$ (top row), $\phi = 3\pi/4$ (centre-top row), $\phi = 5\pi/4$ (centre-bottom row) and $\phi = 7\pi/4$ (bottom row). The left and right columns correspond to the first and third copula examples, respectively.

C.5 Evaluating the effect of sample size

Figures ER-FF illustrate the effect of the sample size T on the boundary set estimates.




Figure ER: Boundary set estimates as $t = 1$ across $T \in \{5,000, 10,000, 25,000\}$ for the first copula example.




Figure ES: Boundary set estimates as $t = T/2$ across $T \in \{5,000, 10,000, 25,000\}$ for the first copula example.




Figure ET: Boundary set estimates at $t = T$ across $T \in \{5,000, 10,000, 25,000\}$ for the first copula example.




Figure EU: Boundary set estimates as $t = 1$ across $T \in \{5,000, 10,000, 25,000\}$ for the second copula example.




Figure EV: Boundary set estimates as $t = T/2$ across $T \in \{5,000, 10,000, 25,000\}$ for the second copula example.




Figure EW: Boundary set estimates at $t = T$ across $T \in \{5,000, 10,000, 25,000\}$ for the second copula example.




Figure EX: Boundary set estimates as $t = 1$ across $T \in \{5,000, 10,000, 25,000\}$ for the third copula example.




Figure EY: Boundary set estimates as $t = T/2$ across $T \in \{5,000, 10,000, 25,000\}$ for the third copula example.




Figure EZ: Boundary set estimates at $t = T$ across $T \in \{5,000, 10,000, 25,000\}$ for the third copula example.




Figure FA: Boundary set estimates as $t = 1$ across $T \in \{5,000, 10,000, 25,000\}$ for the fourth copula example.




Figure FB: Boundary set estimates as $t = T/2$ across $T \in \{5,000, 10,000, 25,000\}$ for the fourth copula example.




Figure FC: Boundary set estimates at $t = T$ across $T \in \{5,000, 10,000, 25,000\}$ for the fourth copula example.




Figure FD: Boundary set estimates as $t = 1$ across $T \in \{5,000, 10,000, 25,000\}$ for the fifth copula example.




Figure FE: Boundary set estimates as $t = T/2$ across $T \in \{5,000, 10,000, 25,000\}$ for the fifth copula example.




Figure FF: Boundary set estimates at $t = T$ across $T \in \{5,000, 10,000, 25,000\}$ for the fifth copula example.

D Additional case study figures and tables

Figures FG and FH illustrate ACF plots for Q_t and Q_t^2 , respectively, suggesting an independence assumption is unreasonable for each of the time series.

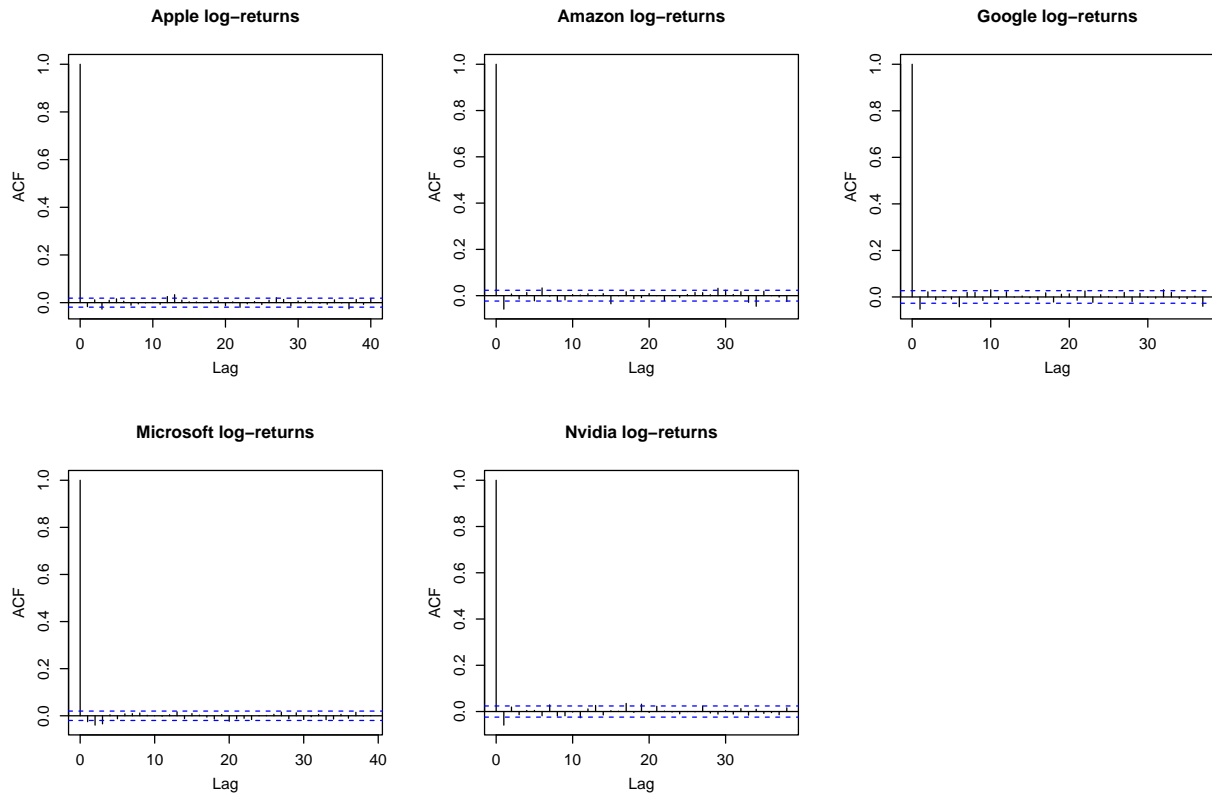


Figure FG: Log-returns ACF plots for each stock.

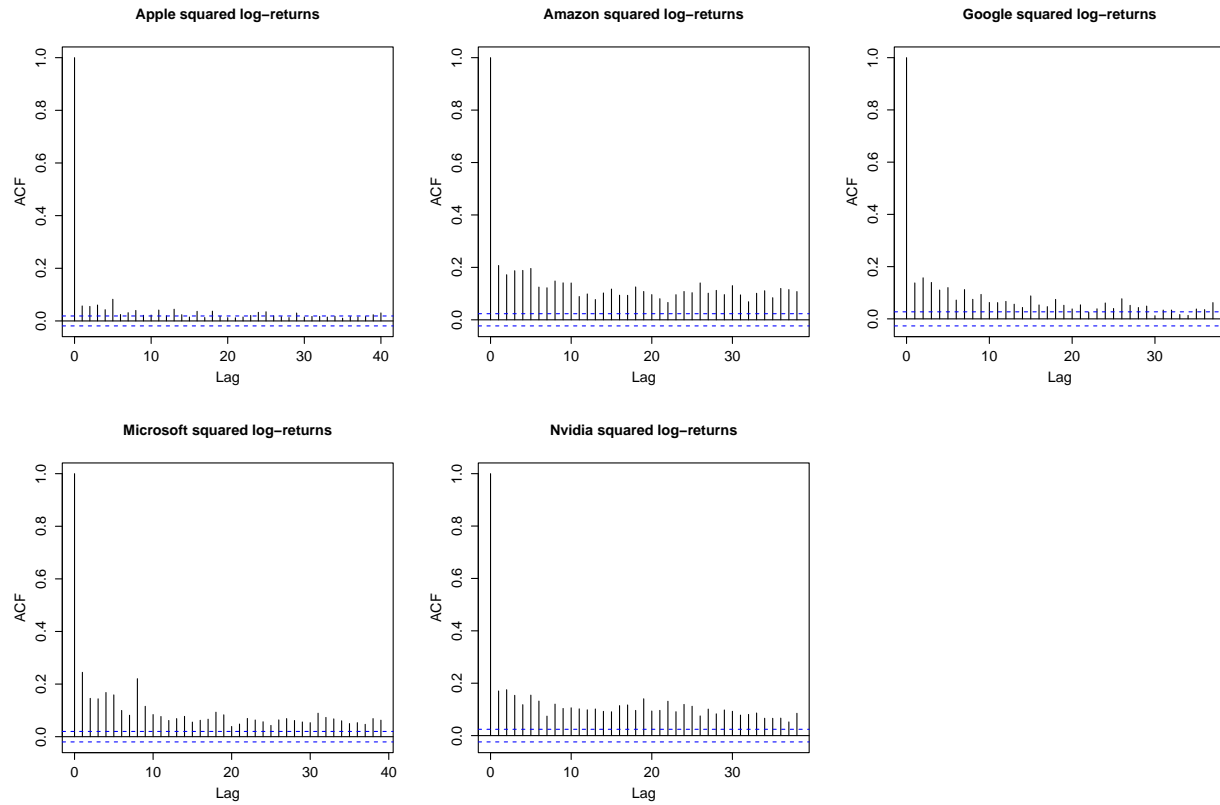


Figure FH: Squared log-returns ACF plots for each stock.

Figures FI and FJ illustrate ACF plots for the filtered and squared filtered log-returns, $\hat{\varepsilon}_t$ and $\hat{\varepsilon}_t^2$, respectively. These plots indicate an independence assumption is reasonable for the residual process.

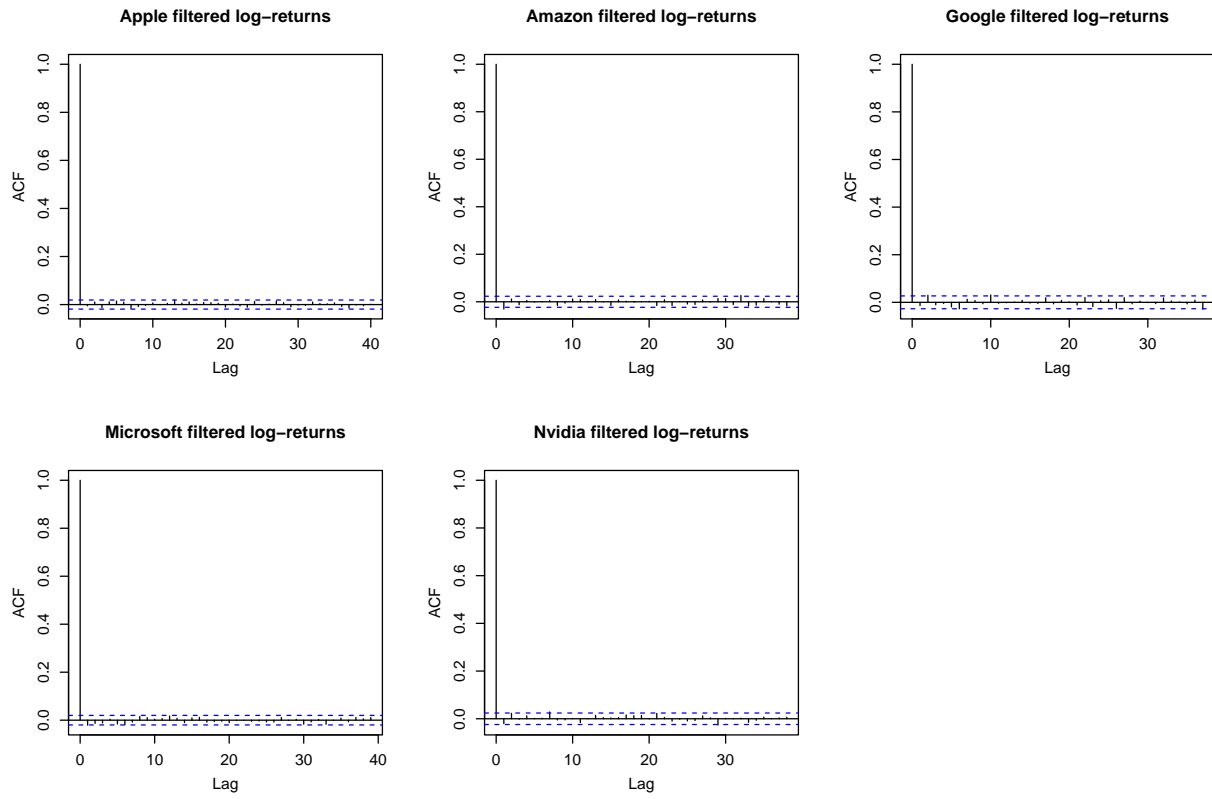


Figure FI: Filtered log-returns ACF plots for each stock.

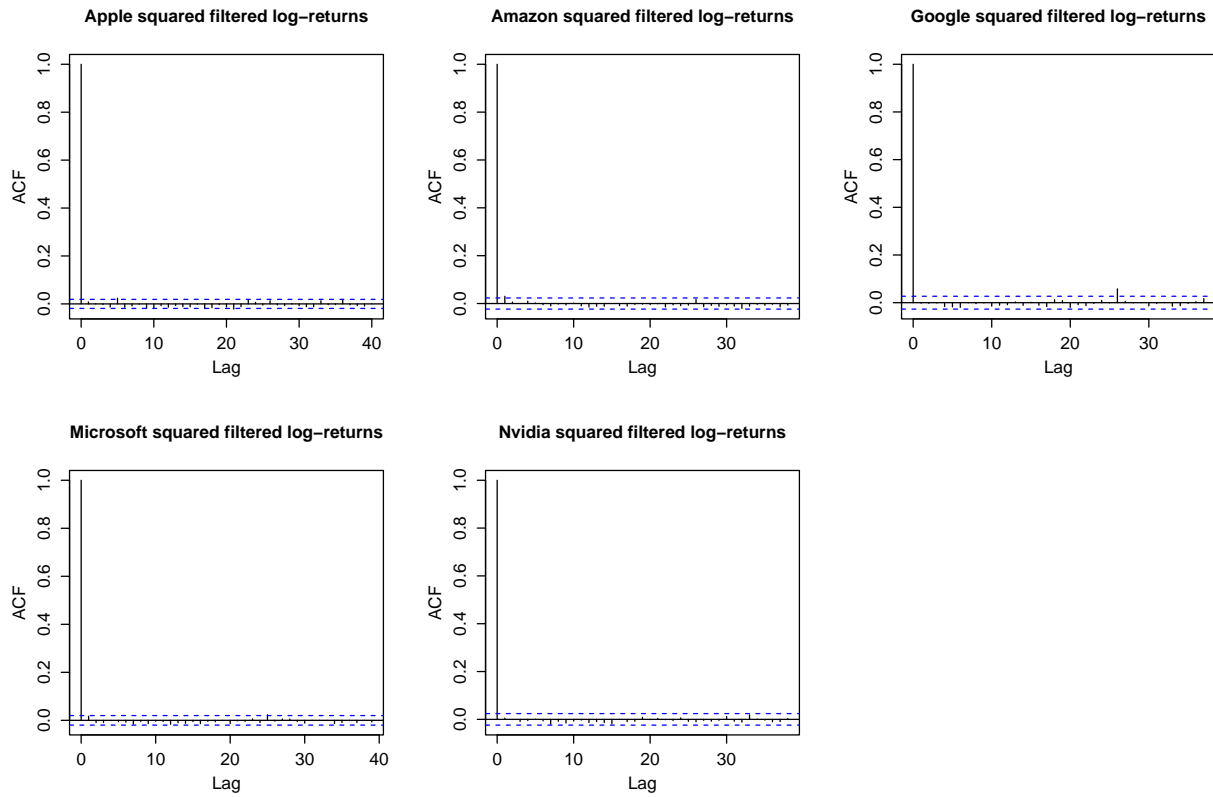


Figure FJ: Squared filtered log-returns ACF plots for each stock.

Figures FK illustrate the individual time series for the filtered log-returns. These plots display no obvious periodicity or long term trends, suggesting a stationarity assumption is reasonable for the filtered data.

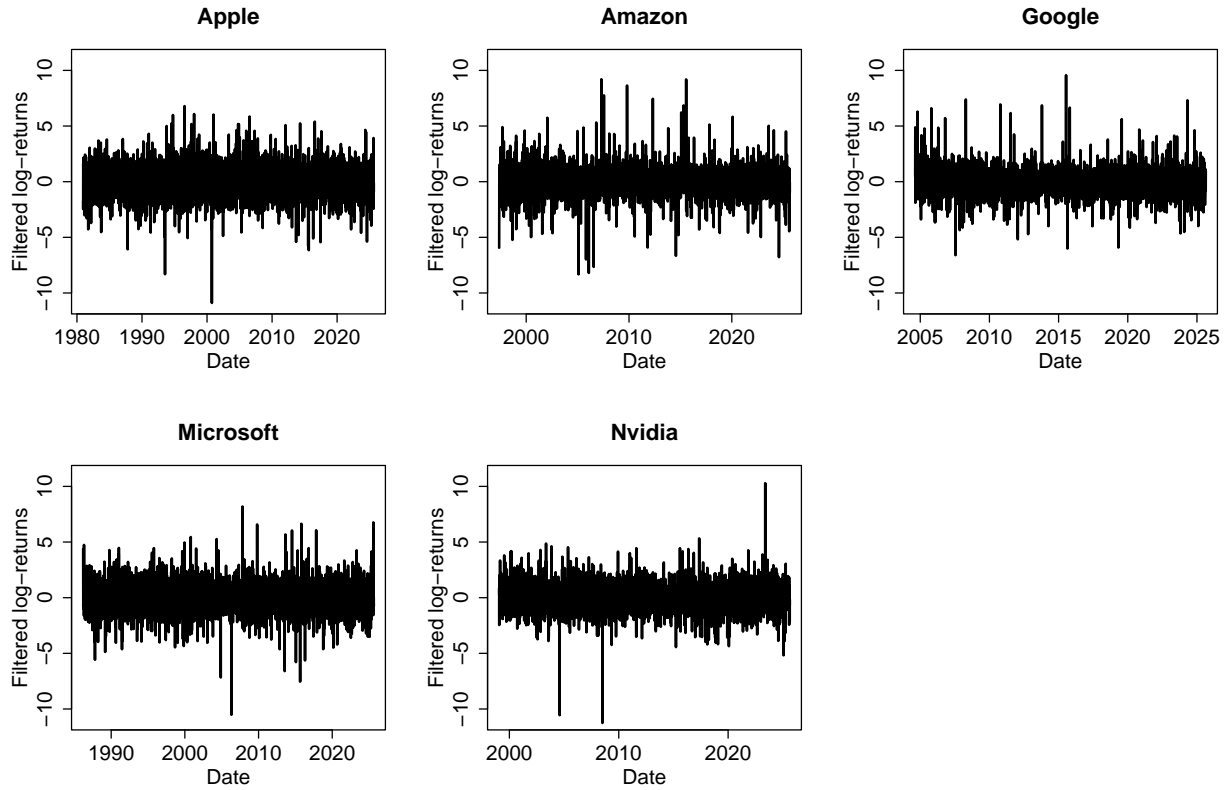


Figure FK: Filtered log-returns time series for each stock.

Figures FL and FO give the QQ plots obtained from the GPD model fits on the lower and upper tails of the filtered log-returns time series. Each plot illustrates model quantile estimates, alongside 95% parametric bootstrapped confidence intervals. One can observe generally good agreement between the observed and estimated quantiles, with only mild disagreements at the most extreme observations for some stocks (e.g., Nvidia), suggesting the proposed marginal model is reasonable for capturing both tails. Note that we fit the GPD model to the Apple data multiple times, due to the fact observation lengths and zero movement days vary significantly between different pairs of stocks, meaning the Apple time series changes for each pairing.

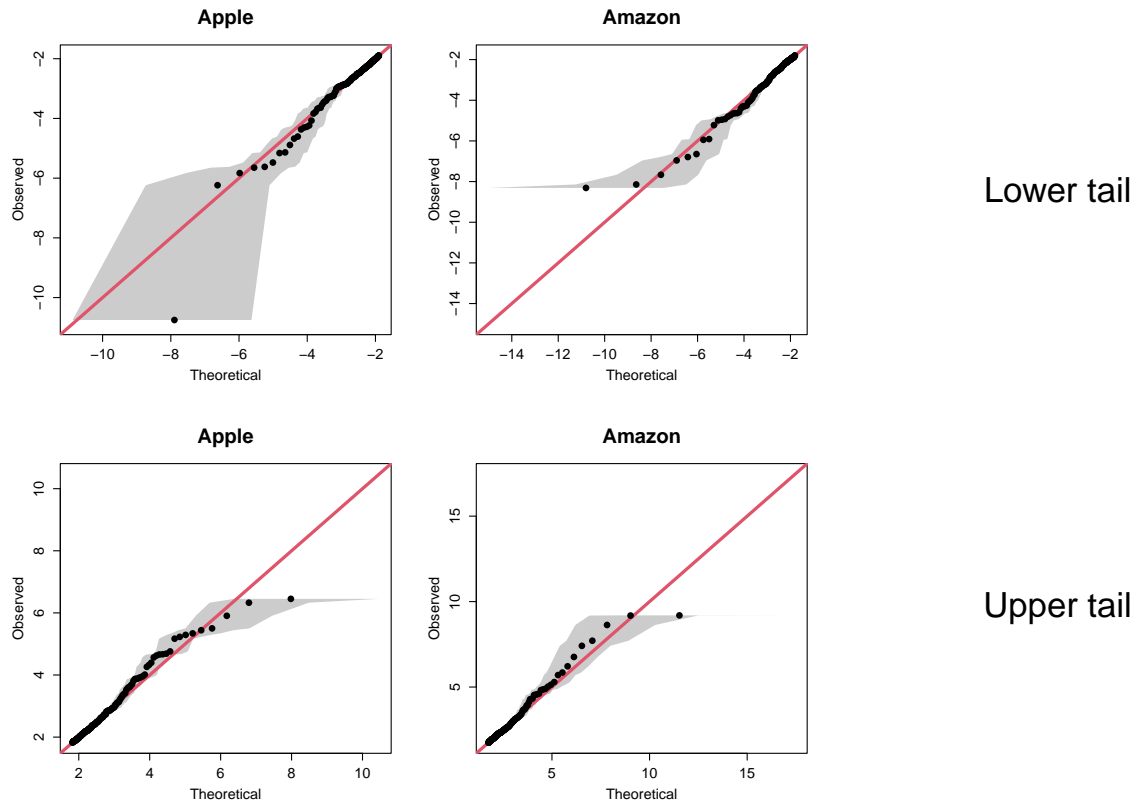


Figure FL: GPD QQ plots for the lower (top row) and upper (bottom row) tails of the Apple and Amazon series.




Figure FM: GPD QQ plots for the lower (top row) and upper (bottom row) tails of the Apple and Google series.




Figure FN: GPD QQ plots for the lower (top row) and upper (bottom row) tails of the Apple and Microsoft series.




Figure FO: GPD QQ plots for the lower (top row) and upper (bottom row) tails of the Apple and Nvidia series.

Figures FP and FS illustrate rolling window Laplace distribution parameter estimates for each pair of stocks post marginal transformation. In all cases, the rolling window estimates remain close to the ‘true’ standard Laplace parameters, which are also shown in the panels.




Figure FP: Rolling window Laplace parameter estimates for the transformed Apple and Amazon series.




Figure FQ: Rolling window Laplace parameter estimates for the transformed Apple and Google series.




Figure FR: Rolling window Laplace parameter estimates for the transformed Apple and Microsoft series.




Figure FS: Rolling window Laplace parameter estimates for the transformed Apple and Nvidia series.

Figure FT illustrates the estimated bivariate extremograms for the filtered volatility series of each pair of stocks. The extremograms are defined using 0.95 empirical quantiles from each series. These plots suggest that assuming extremal independence between the filtered series at different time lags is reasonable.




Figure FT: Bivariate extremograms for the squared residual series of each pair of stocks.

Figure FU illustrate the PP plots associated with the truncated gamma model fits across all pairs. Furthermore, Figure FV illustrates the differences of the PP plot model and theoretical probabilities over the corresponding rank.



Figure FU: Truncated gamma PP plot diagnostic for each pair of stocks. The red lines denote equity, and the grey shaded regions denote 95% pointwise confidence intervals.




Figure FV: Differences between PP plot model and theoretical probabilities for each pair of stocks. The red lines denote equity, and the grey shaded regions denote 95% pointwise confidence intervals.

Figure FW illustrates the time-varying threshold quantile sets for each pair of filtered log-returns. Such sets can be seen at the ‘threshold level’ above which we fit the truncated gamma model.




Figure FW: Threshold quantile set estimates $\widehat{\mathcal{R}}_t^\tau$ at $\tau = 0.8$ over time for each pair of stocks. The colour scale is used to illustrate the variation over time, with the blue and orange sets corresponding to the start and end of the observation periods, respectively.

Figure FX illustrates return level set estimates at $p = 0.999$ for significant dates on the standard Laplace scale. Note that we observe only subtle differences between the estimated sets in each case.




Figure FX: Return level set estimates on the standard Laplace scale at $p = 0.999$ across all pairs for recent dates of significance in the financial context. The legend is the same as in Figure 10.

Figure FY illustrates the return level set probabilities diagnostic for each pair of stocks. One can observe good agreement between the tuples in every case.




Figure FY: Return level set diagnostic for each pair of stocks. The red lines denote equity.

The time-varying η estimates for each pair in quadrants 2 and 4 are shown in Figure FZ. As in Section 5.5, we observe a range of complex dependence trends across the different pairs.




Figure GA: Downside and upside VaR and CoVaR estimates of $X_{2,t}$ for all pairs.




Figure FZ: The coefficient of tail dependence η_t over time for quadrants 2 (top row) and 4 (bottom row). The black dotted line denotes the upper bound for η_t .

Figure GA illustrates the upside and downside VaR and CoVaR estimates of $X_{2,t}$ for all pairs. These plots illustrate the impact of the stress events from $X_{1,t}$; if each pair were independent, then the VaR and CoVaR estimates of $X_{2,t}$ would be identical.

Figure GB illustrates the upside and downside VaR and CoVaR estimates for all pairs on the Laplace scale. These plots illustrate the complex evolution of the extremal dependence structure over time.




Figure GB: Downside and upside VaR and CoVaR estimates for all pairs on the Laplace scale, with the 'stress' events defined by the Apple stock in each case.