Risk analysis of cumulative intraday return curves

Piotr Kokoszka[*], Hong Miao[†], Stilian Stoev[‡], Ben Zheng[§]

May 10, 2018

Abstract

Motivated by the risk inherent in intraday investing, we propose several ways of quantifying extremal behavior of a time series of curves. A curve can be extreme if it has shape and/or magnitude much different than the bulk of observed curves. Our approach is at the nexus of Functional Data Analysis and Extreme Value Theory. The risk measures we propose allow us to assess probabilities of observing extreme curves not seen in a historical record. These measures complement risk measures based on point-to-point returns, but have different interpretation and information content. Using our approach, we study how the financial crisis of 2008 impacted the extreme behavior of intraday cumulative return curves. We discover different impacts on shares in important sectors of the US economy. The information our analysis provides is in some cases different from the conclusions based on the extreme value analysis of daily closing price returns.

Key words: Cumulative intraday returns, Extremes, Functional data, Risk measures.

1 Introduction

A fundamental concept of quantitative finance is the point–to–point return. For example, daily returns are defined as

$$r_n = 100(\ln P_n - \ln P_{n-1}) \approx 100\frac{P_n - P_{n-1}}{P_{n-1}}, \tag{1.1}$$

where P_n is the closing price on trading day n. The unit of time can be day, month, or year, with the classical finance theory initiated by (Markowitz, 1959), (Sharpe, 1964), (Lintner, 1965), and (Black, 1972) pertaining to annual returns. This paper is concerned with daily curves of the cumulative intraday returns defined as follows.

[*]Department of Statistics, Colorado State University, Fort Collins, CO 80523
[†]Department of Finance and Real Estate, Colorado State University, Fort Collins, CO 80523
[‡]Department of Statistics, University of Michigan, Ann Arbor, MI 48109
[§]Department of Statistics, Colorado State University, Fort Collins, CO 80523

DEFINITION 1.1 Suppose $P_n(t)$, $n = 1, \ldots, N$, is the price of a financial asset at time t on trading day n. The functions

$$R_n(t) = 100[\ln P_n(t) - \ln P_n(t_0)], \qquad n = 1, \ldots, N,$$

are called the cumulative intraday returns (CIDRs).

In practice, the CIDRs can be computed only at discrete time points, but the resolution can be very fine for frequently traded assets and indexes. In this research, we use one-minute resolution with t_0 corresponding to the first minute after the market opening; $P_n(t)$ is the average of the maximum and minimum price within the minute. For illustration, five consecutive CIDR curves are displayed in Figure 1.1. These curves show how the return accumulates throughout a trading day. Since $R_n(t) \approx 100(P_n(t) - P_n(t))/P_n(t_0)$, with $P_n(t_0)$ being a constant for a given day n, the price and the CIDR curves look very similar. However, the CIDR curves are scale and level normalized, so they form a stationary sequence in a function space, (Horváth *et al.*, 2014). We view CIDR curves as a time series of functions, and focus is on the analysis of their shapes. For this reason, we adopt the general framework of functional data analysis. While a substantial body of work on various aspects of time series of functions has accumulated, as we discuss below, we are not aware of any work on the *extremes* of such time series. This paper is an attempt to provide a possible definition of extreme curves and study their properties in the context of CIDRs. Extreme value analysis of point–to-point returns occupies a prominent place in econometric research due to its direct relation to risk analysis, but no work has been done on the extreme shapes of the CIDR *curves*.

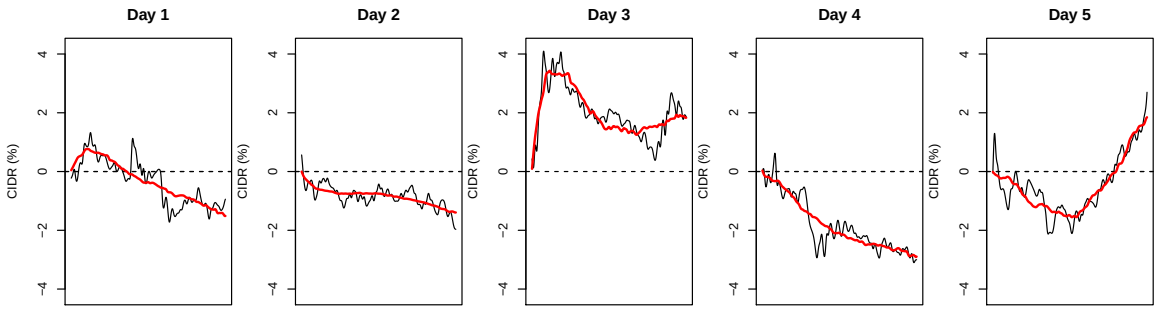

Figure 1.1: Five Consecutive CIDRs (sector ETF XLF). The noisy curves are the raw CIDRs, the superimposed thick curves are the corresponding CIDRs reconstructed from the first three FPCs. Details are explained in Section 2.

The CIDR curves describe how the return on an investment evolves with time over a relatively short period. The information contained in them is different than the information provided by minute-by-minute or trade-by-trade returns, and may be more relevant

in situations where positions are held for a period of few hours within a trading day. Risk measures computed for point–to–point returns may be not sufficient for evaluating risks for portfolios held by such trades. For example, valuations of certain options are impacted by the shape the price curve may take over some interval, and knowing the probabilities of prespecified unusual shapes may be useful. The type of risk we study is different from risk related to micro–structure noise. The latter is relevant to high frequency trades executed within seconds. The risk inherent in the shapes of cumulative returns is relevant to positions held at least for several hours. Micro-structure noise is relevant to professional traders, whereas smooth within day price evolution is relevant to a broader, and growing, group of investors owning, for example, Exchange Traded Funds, which can be bought or sold within a trading day. This time horizon is in turn shorter than that of longer term investors, holding, for example, positions in Mutual Funds.

Our objective is to develop a possible general framework, we do not develop any asymptotic theory at this stage, but rather focus on illustrating the concepts by means of various graphical displays.

The paper is organized as follows. In the remainder of the introduction, we cite a few selected monographs and papers related to the themes of our work: functional data analysis and risk analysis of returns. These fields of research are large, so our small review is necessarily selective and biased toward work we are most familiar with. Section 2 presents statistical preliminaries, including the most relevant methodology of functional data analysis and extreme value theory. Using these tools, we conduct a risk analysis of several aspects CIDRs in Sections 3. We focus on the impact of the financial crisis of 2008 on the extremal behavior of the CIDR curves. Large tables and graphical displays are collected in Appendix D, so that they do not interrupt the narrative.

General introductions to functional data analysis (FDA) are given in (Ramsay and Silverman, 2005) and (Kokoszka and Reimherr, 2017). Methodology based on functional principal components, which we use extensively here, is the focus of (Horváth and Kokoszka, 2012). Early results on linear functional time series, including their prediction are presented in (Bosq, 2000). Ideas of FDA have been applied to inference for CIDRs by (Kokoszka and Reimherr, 2013), (Lucca and Moench, 2015), (Kokoszka *et al.*, 2015), (Zhang, 2016), (Kokoszka *et al.*, 2017) and (Kokoszka *et al.*, 2017b).

Extreme value theory (EVT) with a view toward risk analysis is presented in (Embrechts *et al.*, 1997) and (McNeil *et al.*, 2005), monographs which define and study the value-at-risk (VaR) and the expected shortfall (ES) (we review these concepts in Section 2). Extreme behavior of *point–to–point* returns has been studied in many papers. For example, (Bensalah, 2000) applied univariate EVT techniques to a series of daily exchange rates of Canadian/U.S. dollars over a 5-year period (1995-2000), (Gençay *et al.*, 2003) compared the performance of the extreme value theory in VaR calculations to other well-known modeling techniques and concluded that the EVT provides an approach which is robust and easy to implement. (Gilli and Këllezi, 2006) used the extreme value theory to compute VaR and ES and related confidence intervals of several major stock market

indices for the period 1960–2004. (Kourouma *et al.*, 2011) investigated VaR and ES for several market indexes during the 2008 financial crisis. Other papers discussing the tail behavior of financial series include (Danielsson and de Vries, 2000), (Neftci, 2000), (McNeil and Frey, 2000), (Gençay and Selçuk, 2004), (Brooks *et al.*, 2005) and (Brodin and Klüppelberg, 2008). (Rocco, 2014) provides a critical review of the theoretical underpinnings of extreme value theory and a survey of some major applications of extreme value theory to finance. Intraday VaR was initially discussed by (Giot, 2005) who estimates a conditional parametric VaR in which intraday volatility is modeled using different specifications of GARCH, and RiskMetrics models on 15- and 30-minute returns of three stocks traded on the New York Stock Exchange. More recently, (Dionne *et al.*, 2009) proposed an approach to estimate intraday Value at Risk (IVaR) using tick-by-tick data. Value-at-risk or Conditional Value-at-risk using intraday price data contains richer information than the traditional daily data based measures and are proper for shorter term portfolios. Our analysis is different, it focuses on extreme shapes of the intraday return curves.

2 Statistical preliminaries

This section introduces the requisite statistical background in FDA and EVT. Readers familiar with these fields may scan this section to become acquainted with the notation we use.

2.1 Background on Functional Data Analysis

Square integrable functions We assume that all functions are defined on a compact interval, which, for simplicity, we take to be the unit interval $[0, 1]$. A function f is said to be square integrable if $\int_0^1 f^2(t)dt < \infty$. We denote the set of all the square integrable functions by L^2. It is a Hilbert space with the inner product

$$\langle f, g\rangle = \int_0^1 f(t)g(t)dt, \quad f, g \in L^2.$$

We say that functions f and g are orthogonal if $\langle f, g\rangle = 0$. The inner product also allows us to define the *norm* as

$$\|f\| = \sqrt{\langle f, f\rangle} = \left\{ \int_0^1 f^2(t)dt \right\}^{1/2}.$$

Basis expansions play an important role in FDA. We say that functions $\{e_1, e_2, \dots\}$ form a basis in L^2 if every function $f \in L^2$ admits a unique expansion $f(t) = \sum_{j=1}^{\infty} a_j e_j(t)$. The basis is orthonormal if $\langle e_j, e_{j'}\rangle = 0$ whenever $j \neq j'$ and $\|e_j\| = 1$. In that case, Parseval's identity yields

$$\|f\| = \left\{ \sum_{j=0}^{\infty} a_j^2 \right\}^{1/2}, \quad a_j = \langle f, e_j\rangle .$$

In the statistical FDA framework, the observations are functions. In our setting, we treat the CIDR curve R_n as a realization of a random function R, a random process showing how the return stochastically evolves within a trading day. This means that the norm $\|R\|$, is a random variable. If $E\|R\|^2 < \infty$, we say the random function R is *square integrable*.

The Karhunen-Loéve decomposition Consider a random curve $X = \{X(t),\ t \in [0,1]\}$ in the space L^2. We define the mean and covariance functions by $\mu(t) = EX(t)$ and $c(t,s) = E[(X(t) - \mu(t))(X(s) - \mu(s))]$. The Karhunen-Loéve decomposition states that every square integrable function X can be represented as

$$X(t) = \mu(t) + \sum_{j=1}^{\infty} \xi_j v_j(t). \tag{2.1}$$

The v_j are the eigenfunctions of c, i.e. they are the solutions to the equation

$$\int c(t,s)v(s)ds = \lambda v(t),$$

where the corresponding λ_j are called the eigenvalues of X (or c). They are usually arranged in non-increasing order: $\lambda_1 \geq \lambda_2 \geq \dots$. The random variables ξ_j, called the *scores*, are given by

$$\xi_j = \langle X - \mu, v_j \rangle = \int (X(t) - \mu(t))v_j(t)dt.$$

Decomposition (2.1) is optimal in a sense that X can be well approximated using only very few initial v_j, in many applications 2 or 3 of them. The functions v_j form an orthonormal basis and are also called the *Functional Principal Components* (FPCs) of X. It can be shown that

$$E\xi_j = 0, \quad E\xi_j^2 = \lambda_j, \quad \mathrm{Cov}(\xi_j, \xi_k) = 0, \ \text{ if } j \neq k$$

and

$$E\|X - \mu\|^2 = \sum_{j=1}^{\infty} \lambda_j. \tag{2.2}$$

The interpretation of the above formula is that λ_j is the variance of the random function X in the principal direction v_j. The sum of these variances is the total variance of X. Relation (2.2) is thus the decomposition of variance corresponding to (2.1).

The above expansions have sample analogs. If $X_1, X_2, \dots, X_N$ is a sample of functions, we define

$$\hat{c}(t,s) = \frac{1}{N}\sum_{i=1}^{N} \big(X_i(t) - \hat{\mu}(t)\big)\big(X_i(s) - \hat{\mu}(s)\big), \quad \hat{\mu}(t) = \frac{1}{N}\sum_{i=1}^{N} X_i(t)$$

and

$$\int \hat{c}(t,s)\hat{v}_j(s)ds = \hat{\lambda}_j\hat{v}_j(t). \tag{2.3}$$

If the functions X_i form a stationary, weakly dependent time series in L^2, the approximation

$$X_i(t) \approx \hat{\mu}(t) + \sum_{j=1}^{p} \hat{\xi}_{ij}\hat{v}_j(t), \quad \hat{\xi}_{ij} = \langle X_i, \hat{v}_j \rangle,$$

is generally very accurate for small p. In Section 3, we will see that for the CIDR curves $p = 2$ or $p = 3$ is sufficient. The above sample approximation allows us to reduce the analysis of statistical properties of curves to the analysis of statistical properties of vectors of reasonably low dimension.

2.2 Background on Extreme Value Theory

Peaks-over-Threshold methodology The peaks-over-threshold (PoT) analysis of univariate time series has been widely used since the seminal paper of (Davison and Smith, 1990). They advocated the use of the asymptotically motivated Generalized Pareto distribution as a model for the distribution of exceedances over a certain high threshold. Let $Y_1, Y_2, \ldots, Y_N$ be a sequence of *scalar* observations with a common marginal distribution F, $F(y) = P(Y \leq y)$. We are interested in estimating the *conditional excess distribution function*

$$F_u(y) = P(Y \leq u + y \mid Y > u) = \frac{F(u+y) - F(u)}{1 - F(u)}, \qquad y > 0,$$

over the threshold u. For large enough u, under mild conditions,

$$F_u(y) \approx G_{\gamma,\sigma_u}(y) := 1 - \left(1 + \frac{\gamma}{\sigma_u} y\right)_+^{-1/\gamma}, \tag{2.4}$$

where γ is called the *shape parameter* and $1/\gamma$ the *tail index*. The distribution G_{γ,σ_u} is called the Generalized Pareto Distribution (GPD). Relation (2.4) can be equivalently stated, for $x > u$, as

$$P(Y > x \mid Y > u) \approx 1 - G_{\gamma,\sigma_u}(x-u) = \left[1 + \gamma\left(\frac{x-u}{\sigma_u}\right)\right]_+^{-1/\gamma}.$$

It follows that

$$P(Y > x) \approx P(Y > u)\left[1 + \gamma\left(\frac{x-u}{\sigma_u}\right)\right]_+^{-1/\gamma}, \qquad u \to \infty.$$

The tail index $1/\gamma$ quantifies the heaviness of the tail. Namely, if $\gamma > 0$, then the asymptotic distribution of the excesses is unbounded and has heavy, Pareto-like tail,

$P(Y > x) \sim x^{-1/\gamma}, x \to \infty$. The larger the index γ, the heavier the tail. The case $\gamma < 0$ corresponds to excesses with finite upper bound. Observe that the condition $1 + \gamma(x-u)/\sigma_u \geq 0$ means that $x \leq u + \sigma_u/|\gamma| < \infty$. Finally, the boundary case $\gamma = 0$ corresponds to the exponential distribution since

$$\lim_{\gamma \to 0} \left[1 + \gamma\left(\frac{x-u}{\sigma_u}\right)\right]_+^{-1/\gamma} = \exp\left\{-\frac{x-u}{\sigma_u}\right\}, \quad x \geq u.$$

The PoT method involves threshold selection, namely, how large should u be so that the approximation in (2.4) is accurate. Usually, one can assess the stability of parameter estimates, based on fitting of models across a range of threshold values. Alternatively, this can be done by drawing a mean residual life plot. We produce the mean residual life plot which estimates the function $u \mapsto \mathbb{E}\big(Y - u \mid Y > u\big)$. If Y follows a GPD, then the above function should be linear in u. More specifically, if u_0 is a threshold, $\mathbb{E}\big(Y - u_0 \mid Y > u_0\big) = \sigma_{u_0}/(1-\gamma)$ exists, if $\gamma < 1$. Then for any $u > u_0$,

$$\mathbb{E}\big(Y - u \mid Y > u\big) = \frac{\sigma_u}{1-\gamma} = \frac{\sigma_{u_0} + \gamma(u - u_0)}{1-\gamma}. \tag{2.5}$$

When plotting empirical estimates of the sample means excesses against a range of thresholds, the threshold is chosen to be the lowest level where all the higher threshold based sample mean excesses are consistent with a straight line, once the sample uncertainty is accounted for.

Other traditional ways of choosing threshold involve some values pre-determined by physical considerations or some simple fixed quantile rule, like the upper 10% rule.

Parameter Estimation Having determined a threshold, the parameters can be estimated by the maximum likelihood approach. Suppose that the total number of observations is N, and N_u of them exceed the threshold u. We denote them by the order statistics $x_{(N-N_u+1)}, \ldots, x_{(N)}$. Then the log-likelihood derived from (2.4) is

$$l(\sigma_u, \gamma) = -N_u \log \sigma_u - \left(1 + \frac{1}{\gamma}\right) \sum_{i=1}^{N_u} \log\left(1 + \frac{x_{(N-i+1)} - u}{\sigma_u}\gamma\right), \tag{2.6}$$

provided $(1 + \sigma^{-1}(x_{(N-i+1)} - u)\gamma) > 0$, for $i = 1, \ldots, N_u$. Closed form maximization of the log-likelihood is not possible, so numerical techniques are utilized. Standard errors and confidence intervals for the corresponding parameters are obtained from the likelihood theory. The required asymptotic normality of the maximum likelihood estimators holds for $\gamma > -1/2$. In our context, γ is close to zero.

In practice, the estimation of the parameters γ and σ can be done using the function `gpd.fit` in the R package `ismev`. Estimating $P(Y > u)$ by N_u/N, with the ML estimates

$\hat{\gamma}$ and $\hat{\sigma}_u$ obtained from the above procedure leads to

$$\hat{F}(x) \approx 1 - \frac{N_u}{n}\left[1 + \hat{\gamma}\left(\frac{x-u}{\hat{\sigma}_u}\right)\right]_+^{-1/\hat{\gamma}} \tag{2.7}$$

for any $x > u$ such that $1 + \hat{\gamma}(x-u)/\hat{\sigma}_u > 0$.

VaR, return level, and expected shortfall Value-at-Risk (VaR) is one of the most frequently used risk measures in finance. It involves extreme quantile estimation. Suppose a random variable Y with continuous distribution function F models losses or negative returns on an asset over a certain time horizon. Value-at-Risk at the upper tail α, denoted by VaR_α, is defined as the $100(1-\alpha)$-th quantile of the distribution F,

$$\mathrm{VaR}_\alpha(Y) = F^{-1}(1-\alpha),$$

where F^{-1} is the inverse of the distribution F, and α is typically close to zero. If Y models losses, then $\mathrm{VaR}_\alpha(Y)$ is interpreted as the amount of loss exceeded only $100\alpha\%$ of the time, on the average. A risk measure equivalent to VaR is the *return level*. Let $\alpha = 1/m$, then $L_m = \mathrm{VaR}_{1/m}$ is the *m-period return level*. For instance, it is often more convenient to report return levels on an annual scale. There are about 252 trading days in a year. Taking $m = 252$, we get an upper tail α that approximately equals 0.004 and hence $L_{252} = \mathrm{VaR}_{0.004}$. Thus if Y models a daily loss, then $L_{252}(Y)$ gives us the expected level of loss to be exceeded once a year. The expected shortfall (ES), also called *tail conditional expectation*, is defined as the expected loss that exceeds the VaR_α: $\mathrm{ES}_\alpha = \mathbb{E}(Y \mid Y > \mathrm{VaR}_\alpha)$.

Estimation of the risk measures Denote by $\hat{\sigma}$ and $\hat{\gamma}$ the estimated GPD parameters. To estimate the VaR, we need to estimate $F^{-1}(1-\alpha)$. This can be done by solving equation (2.7) for x. Then VaR_α can be estimated as

$$\widehat{\mathrm{VaR}}_\alpha = x_\alpha = \begin{cases} u + \frac{\hat{\sigma}}{\hat{\gamma}}\left[\left(\frac{N_u}{N\alpha}\right)^{\hat{\gamma}} - 1\right] & \text{if} \quad \hat{\gamma} \neq 0 \\ \\ u + \hat{\sigma}\log\left(\frac{N_u}{N\alpha}\right) & \text{if} \quad \hat{\gamma} = 0 \end{cases}$$

To estimate the expected shortfall ES_α, we use the relation

$$\widehat{\mathrm{ES}}_\alpha = \widehat{\mathrm{VaR}}_\alpha + \mathbb{E}\big(Y - \widehat{\mathrm{VaR}}_\alpha \mid Y > \widehat{\mathrm{VaR}}_\alpha\big),$$

where the second term is the expected exceedance over the threshold $\widehat{\mathrm{VaR}}_\alpha$. Since the *mean excess function* for the GPD with $\gamma < 1$ is given in equation (2.5), it is easy to show that

$$\widehat{\mathrm{ES}}_\alpha = \widehat{\mathrm{VaR}}_\alpha + \frac{\hat{\sigma} + \hat{\gamma}\left(\widehat{\mathrm{VaR}}_\alpha - u\right)}{1-\hat{\gamma}} = \frac{\widehat{\mathrm{VaR}}_\alpha + \hat{\sigma} - \hat{\gamma}u}{1-\hat{\gamma}}.$$

Large sample standard errors, leading to confidence intervals, for VaR_α and ES_α can be derived using the delta method, see Appendix A.

3 Risk analysis of CIDR curves

We now turn to study extremes of the CIDR curves defined in the Introduction. We are specifically interested in the quantification of their extremal behavior before, during and after the financial crisis of 2008. However, the methodology we propose is applicable in many other contexts. A chief question in any analysis of this type is what an *extreme curve* is. Traditional EVT has been developed for samples of scalar observations. Such observations can be ordered, so the concept of "large" observations is clear. Extreme observations are generally those that lie beyond the range of the sample, i.e. have not been observed yet. EVT is then used to compute measures of risk based on upper tail approximations, like those explained in Section 2. In Section 3.1, we explain what we mean by extreme curves. Then, in Section 3.2, we introduce the data sets to which our approaches are applied. Sections 3.3 and 3.4 include the results of our analysis, which Section 3.5 summarizes.

3.1 Extreme return curves

We assume that the observed CIDR curves R_n defined in Definition 1.1 are realizations of a random curve R; they form a stationary time series in L^2, with the marginal distribution equal to that of R. The Karhunen-Loéve decomposition introduced in Section 2 states that the random curve R can be represented as

$$R(t) = \mu(t) + \sum_{j=1}^{\infty} \xi_j v_j(t). \tag{3.1}$$

The mean function μ is assumed to be zero (the estimates $\hat{\mu}(t)$ are not significantly different from zero), so the norm of R is given by

$$\|R\| = \left\{ \int_0^1 R^2(t)dt \right\}^{1/2} = \left\{ \sum_{j=1}^{\infty} \xi_j^2 \right\}^{1/2}, \tag{3.2}$$

where the ξ_j are the scores, uncorrelated random variables.

The simplest way to define a "large" CIDR curve is to require that the norm $\|R\|$ be large. By (3.2) this is equivalent to some of the scores being large, irrespective of their sign. There are many ways in which the curve R can be large according to this definition. A very wiggly curve that starts from zero and ends at zero will be large, even though the daily point-to-point return will be zero. Such a situation corresponds to many high frequency scores being large, but the most important first score being close to zero. To explain this point, we show in Figure 3.1 the first three FPCs, v_1, v_2, v_3 for SP500. They look almost identical for all other assets. The first FPC, v_1, which generally explain about 85% of the variability of these curves, quantifies the upward (if $\xi_1 > 0$) or downward daily

trend. This is a pattern that is observed on most trading days. Thus for a zero daily point-to-point return, the coefficient of this component must be close to zero. (The estimated FPC have the same periodicity pattern as those of the Brownian motion, but they rise more steeply at the beginning of the trading day, reflecting the abrupt price changes often seen at the opening of the market.)

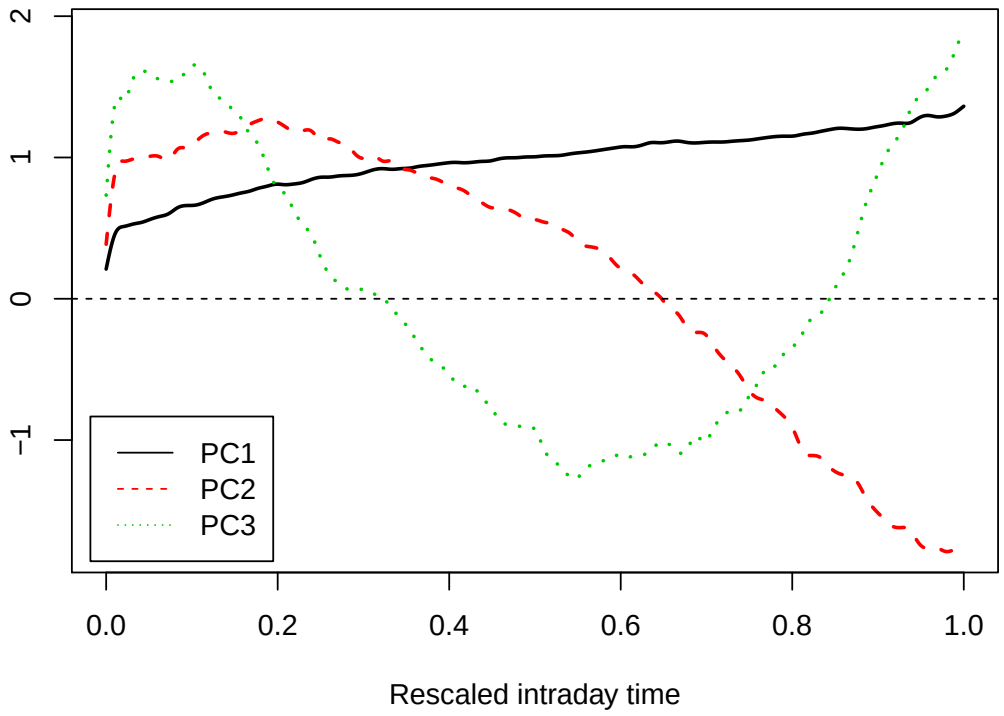

Figure 3.1: The first three FPCs for S&P500 during the period 2006/7/5 - 2011/12/30.

Empirical analysis shows that the shapes of the observed CIDR curves are basically encoded by the first three scores:

$$R_n(t) \approx \xi_{n1} v_1(t) + \xi_{n2} v_2(t) + \xi_{n3} v_3(t).$$

The overwhelming contribution comes from the first and the second FPCs; v_2 quantifies a reversal during a trading day. This is illustrated in Figure 3.2, which shows the plots of $A_n^{(p)} = \{\sum_{j=1}^{p} \xi_{nj}^2\}^{1/2}$, $p = 2, 3, 4$, against $\|R_n\|$ for a selected asset and 50 trading days. The graphs for other assets and periods are similar. We see that $p = 4$ gives basically the same norm as $p = \infty$, and except for "very small" CIDR curves, using $p = 3$ is sufficient. For days with large market swings, using $p = 2$ is sufficient to describe the intraday market action. Thus, in addition to the norm, it is useful to study the extremal behavior of the sequences $\{\xi_{n1}\}$ and $\{\xi_{n2}\}$. Extremes of these sequences will define extreme curves.

In most cases, $\|R\|$ will be large if ξ_1^2 is large, but it does not have to be always the case. The norm $\|R\|$ can be viewed as a new measure of the daily volatility of an asset, but to get a fuller picture of extreme behavior, the first two scores must be considered as well.

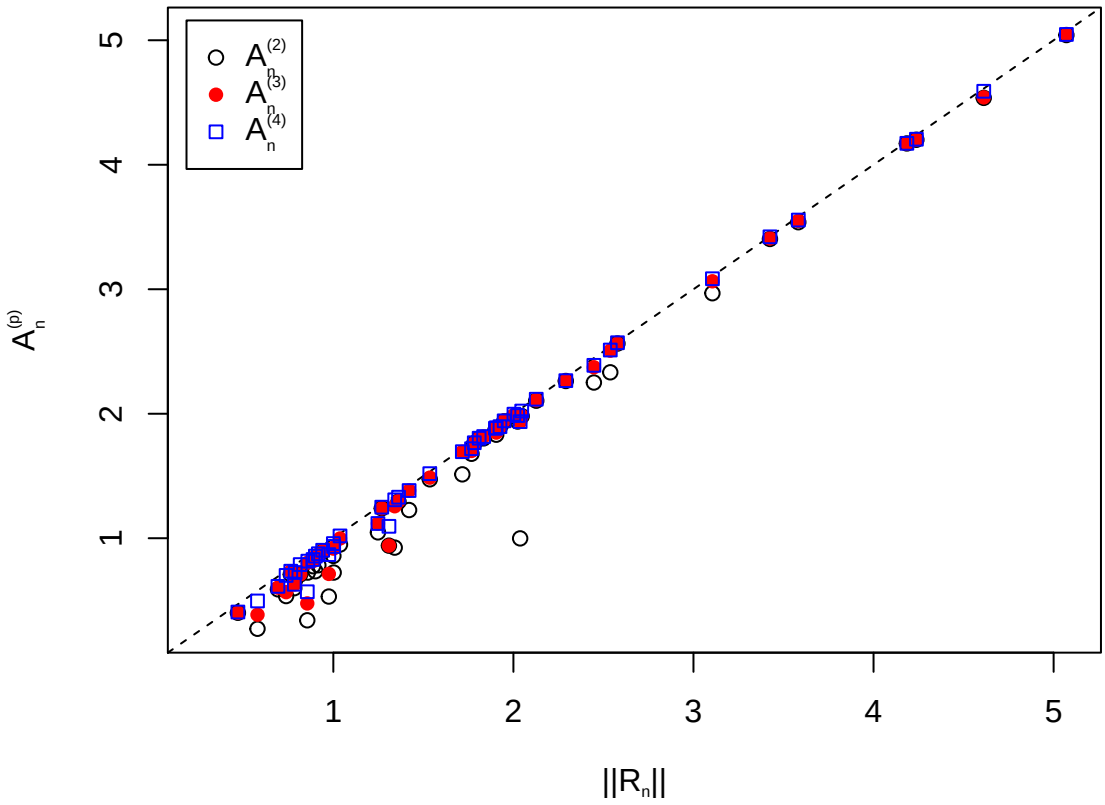

Figure 3.2: The scatter plot of $A_n^{(p)} = \{\sum_{j=1}^{p} \xi_{nj}^2\}^{1/2}$ against $\|R_n\|$ with $p = 2, 3, 4$ for XLF during the period 2008/6/2 - 2008/8/12 (50 days).

To summarize, we will consider a CIDR curve R_n on the trading day n as extreme, if either $\|R_n\|$ is extreme or the scores ξ_{n1} or ξ_{n2} are extreme (or both). The scores can be extreme on either the negative or the positive side.

3.2 Data description

Our application focuses on nine Select Sector SPDR ETFs, which are proxies for "sector portfolios". These are Exchange Traded Funds that track the nine S&P 500 sector indexes, they hold individual stocks within the corresponding sectors. For ease of reference, Table 3.1 lists the nine Select Sector SPDR ETFs. The one minute frequency ETF price data are obtained from Quantquote.

Ticker	Sector	Ticker	Sector
XLF	Financials	XLV	Health Care
XLK	Technology	XLI	Industrials
XLY	Consumer Discretionary	XLB	Materials
XLP	Consumer Staples	XLU	Utilities
XLE	Energy		

Table 3.1: Tickers of the assets (Sector ETFs) used in this study.

One of our major objectives is to examine if the extremal behavior of CIDRs for these "portfolios" is "time-varying". We compare risk measures for periods around the financial crisis. This allows us to determine whether the crisis has a significant impact on the extremal behavior of these assets in a sense defined in Section 3.1. We consider four time periods specified in Table 3.2. The inclusion of two time periods after the crisis is justified by our preliminary findings that indicated a difference between the extremal behavior of the cumulative return curves immediately after the generally accepted end of the crisis and at a later time.

Designation	Abbreviation	Time span	Sample size (days)
Before	Bf	07/05/2006 – 09/28/2007	313
During	Dr	10/01/2007 – 02/27/2009	351
After 1	A1	03/02/2009 – 07/30/2010	356
After 2	A2	08/02/2010 – 12/30/2011	358

Table 3.2: Time periods used in this study.

3.3 Extremes of the norm and the scores

We first apply the PoT method to the norm (3.2) computed using the sum of the squares of the first 15 scores. The application involves the following steps: select the threshold u, fit a GPD to the exceedances over u, compute point estimates and standard errors of the Value-at-Risk and expected shortfall, see Section 2. The results are shown in Table D.1. The columns "u" and "N_u", respectively, show the threshold and the number of observations over that threshold. The column "Parameter $\hat{\gamma}$" reports the estimates of the shape parameter with their standard errors. The shape parameter γ for most assets and periods is seen to be not significantly different from zero, implying that in most cases computation of the risk measures using a Gumbel distribution is appropriate. For longer periods of time, the estimates of γ tend to be positive, which may reflect inhomogeneity of the data or higher risk involved in longer time horizons. To compute the risk measures, we use the GPD with the estimated γ. The last four columns in Table D.1 show the estimation results for the VaR and ES of $\|R\|$ with upper tail probability $\alpha = 0.004$ and

$\alpha = 0.0008$. Since there are 252 trading days, $\text{VaR}_{.004}$ and $\text{VaR}_{.0008}$ are, respectively, the 1-year and 5-year return levels of the norm of the CIDR. The following general pattern for the risk measures for the four periods emerges:

$$\text{during} \;\geq\; \text{after1} \;\geq\; \text{after2} \;\approx\; \text{before}. \tag{3.3}$$

The periods are arranged form the most risky to the least risky. We see that, generally speaking, the impact of the crisis on the extreme risks related to intraday investing persists for about a year and a half after the commonly accepted end of the crisis. Sector specific responses to the crisis are summarized in Figure D.1. At both levels of VaR and ES, financials (XLF) and Consumer Discretionary (XLY) stocks are seen to be most impacted by the crisis; they reach high levels of risk during the crisis. The XLF stocks continue to be exposed to high risk even one year after the crisis. The Technology (XLK) and Health care (XLV) stocks are seen to be, relatively, least impacted by the crisis. They reach the highest risk during the crisis, and then the risk decays almost linearly with time and eventually (2 or 3 years after the crisis) gets back to the level before the crisis. Among all the ETFs, we also found that the Consumer Staples (XLP) and Utilities (XLU) are the most dramatically affected by the crisis. They both have a fairly low risk levels before and immediately after the crisis, while during the crisis, their risk level drastically goes up to a very high level. The remaining sectors, including Energy (XLE), Industrials (XLI) and Materials (XLB), share a similar risk pattern before, during and after the crisis, and they can be categorized as moderately impacted by the crisis.

Since the first FPC usually contributes about 85% to the total variability in the CIDRs and has a clear interpretation as a monotonic growth or decline during the trading day, we applied our risk analysis to the magnitude of the first score, $|\xi_1|$. We show the results in Table D.2. Comparing these results to those in Table D.1, we see a similar pattern for the risk measures among those 4 periods around the crisis. This suggests that the monotonic trend usually dominates the extreme risk of the CIDRs. The pattern of such gains or losses throughout the trading day can also be summarized by relations (3.3).

Comparison to closing price returns We also applied the methodology of Section 2.2 to the traditional point-to-point returns, where only the closing prices on a trading day are used to obtain the daily return. To facilitate the comparison, instead of dealing with left tail (loss) and right tail (gain) of the returns separately, we only consider the magnitude of the point-to-point returns (the absolute value). The results in Table D.3 show a different pattern for the risk measures in the four periods than that revealed by Tables D.1 and D.2. For the point-to-point returns, the impact of the financial crisis can be summarized as follows:

$$\text{during} \;\geq\; \text{after2} \;\geq (\text{or} \approx) \text{after1} \;\geq\; \text{before}. \tag{3.4}$$

The only violation of this pattern is XLF, for which "after1" contains more risk than "after2". Also, another important finding is that for most Sector ETFs, in the "during"

period, the standard errors of VaR and ES are huge, c.f. XLE, XLV and XLU. This indicates that GPD might be an inappropriate model to fit the point-to-point returns in these cases. Moreover, this shows us that the proposed *daily volatility* (or norm of CIDRs) reveals different information from that obtained from the traditional point-to-point returns in the sense of extreme value studies. Sector specific responses to the crisis are summarized in Figure D.2. We took out two assets, XLE and XLU, because of their unreasonable standard errors.

Visualization of extremal regions It is of interest to see how the extreme CIDRs might look like. Since in most cases the shape of the CIDRs is sufficiently well explained by the first two FPCs, and in order to keep the graphs informative, we will visualize the extremal regions by focusing on these two scores. We thus consider the approximation

$$R(t) \approx R^{(2)}(t) := \xi_1 v_1(t) + \xi_2 v_2(t).$$

As extreme, we consider $R^{(2)}$ for which $\|R^{(2)}\| = (\xi_1^2+\xi_2^2)^{1/2}$ is large. We consider extremal regions in four quadrants of the (ξ_1, ξ_2) plane, denoted (+, +), (+, -), (-, -), (-, +), which reflect the four possible combinations of the signs of first two scores. We want to see how these regions change over the "before", "during", "after1" and "after2" periods considered above.

Set $\boldsymbol{\xi} = (\xi_1, \xi_2)$. Given a sufficiently large threshold u, one can postulate the following model for the conditional distribution of $\boldsymbol{\xi}$ given $\|\boldsymbol{\xi}\| > u$:

$$\boldsymbol{\xi} | \|\boldsymbol{\xi}\| > u \overset{d}{\approx} Z(u)\mathbf{Y}(u), \tag{3.5}$$

where $Z(u) \overset{d}{\approx} \|\boldsymbol{\xi}\| | \{\|\boldsymbol{\xi}\| > u\}$ and $\mathbf{Y}(u) \overset{d}{\approx} \boldsymbol{\xi}/\|\boldsymbol{\xi}\| | \{\|\boldsymbol{\xi}\| > u\}$ are independent. The random variable $Z(u)$ follows a GPD and $\mathbf{Y}(u)$ follows some probability distribution on the unit sphere. One can fit various models to the data and use them to compute the probabilities $P(\boldsymbol{\xi} \in A)$, where A is outside the ball $B(0, u)$ with a large radius u. This is done using the relation

$$P(\boldsymbol{\xi} \in A) \approx P(\|\boldsymbol{\xi}\| > u) P(Z(u)\mathbf{Y}(u) \in A).$$

The first factor on the right-hand side can be approximated by using standard univariate EVT. The second factor can be calculated by using Monte Carlo methods.

In the Monte Carlo approach, we first simulate realizations of norms from the model for $Z(u)$. Recall that $Z(u)$ follows a GPD $G_u(z)$, which is fitted from the norms $\|\boldsymbol{\xi}_i\| | \{\|\boldsymbol{\xi}_i\| > u\}$, where $\{\boldsymbol{\xi}_i = (\xi_{1,i}, \xi_{2,i}), i = 1, \ldots, N_u\}$, and N_u is the number of observations with norm greater than u. One can easily draw a new sample of norms $\{z_1, \ldots, z_M\}$ from $G_u(z)$, which can be further used to construct new extreme scores. In the second step, we need to simulate data for the angles $\mathbf{Y}(u)$. To this end, we normalize those observed

pairs $\boldsymbol{\xi}_i = (\xi_{1,i}, \xi_{2,i})$ whose norm $\|\boldsymbol{\xi}_i\|$ exceeds u to get a pool of angles for the extreme points. That is,

$$\theta_i := \arctan(\xi_{2,i}/\xi_{1,i}),$$

for $i = 1, \ldots, N_u$. Based on those observed angles given norms exceeding u, there will be two major ways to generate new angles for future simulation. One could naturally apply a nonparametric method by directly resampling angles from $\{\theta_i, i = 1, \ldots, N_u\}$ with replacement. This is just simply a Bootstrap method. Alternatively, we can also employ a kernel density estimator, see Appendix B for the details of this method. From either of these methods, we can simulate a new set of angles $\{\theta_i^*, i = 1, \ldots, M\}$. Together with the previously simulated norms $\{z_1, \ldots, z_M\}$, we can create simulated score pairs via

$$\boldsymbol{\xi}_i^* = (\xi_1^*, \xi_2^*) = (z_i \cos\theta_i^*, z_i \sin\theta_i^*)$$

with $\|\boldsymbol{\xi}_i^*\| > u$ for $i = 1, \ldots, M$. These newly generated scores can be used to construct a set of simulated extreme curves to depict the extremal region of the CIDRs.

To keep the number of figures relatively small, we merely illustrate our methodology using the XLF (financials) data. The extremal regions obtained using bootstrap are shown in Figures D.3–D.6. Specifically, in Figure D.3, we show the extreme region determined by the norms of the CIDRs on the plane of the scores (ξ_1, ξ_2). The inner black circle on each scatter plot represents the one-year return level ($\text{VaR}_{.004}$) for each of the 4 periods and the bigger grey circle represents a 10,000 day return level ($\text{VaR}_{.0001}$) which can be considered very rarely exceeded. This extremely high level is needed for simulation and visualization, we need a finite upper bound to produce the graphs. We can see for each of the 4 periods that there are one or two observations, represented by stars, falling into the extreme region defined by $\text{VaR}_{.004}$ and $\text{VaR}_{.0001}$, which approximately represents a region with an upper probability of 0.4%. On the same plot, the small grey circles in the extreme region represent around 100 simulated points. With these simulated points, in Figure D.4, we show the regions of extreme shapes of the CIDRs for each quadrant and each period corresponding to the extremal regions shown in Figure D.3. We also plot the real extreme CIDRs, represented by thick solid curves, in the extremal region into which they fall. Similarly, in Figures D.5 and D.6, we show the extreme region relative to the one-year expected shortfall ($\text{ES}_{.004}$) on the plane of the scores (ξ_1, ξ_2), and the corresponding extreme shapes. The expected shortfall is a *coherent* risk measure and constrains both the probability of a tail event and the expected loss given a tail event. This is seen in Figure D.5, where the observed extreme points are very close to the circles that represent norms equal to $\text{ES}_{.004}$. Alternatively, as we mentioned above, the procedure of simulating angles for points with norm greater than u can be also done by using the a kernel density estimator (KDE). The corresponding results are shown in Figures D.7–D.10. A major difference between the two methods in depicting the extreme region can be seen by comparing Figure D.3 and Figure D.7, which show the extreme regions of the score pairs using the VaR estimates. With the bootstrap method, we draw samples with

angles that behave fairly similarly to the observed data, while the KDE method leads to a relatively higher chance of drawing some large angles which are actually rarely observed in the existing data. For example, a score pair with score 1 approximately equal to score 2 is almost never observed in practice, because the score of the first FPC explains most (about 85%) of the total variability, which leaves very low chance for score 2 to reach a value as high as score 1. Figure D.4 and Figure D.8 show that the KDE method leads to a more extended extreme region than the bootstrap method. Similar comments apply to the difference between the two methods for the extreme regions based on the ES.

A general conclusion is that while less intuitive, visualizing extreme curves in the space of scores is more informative, especially if the temporal evolution of these regions is to be examined. Using simple bootstrap may lead to more realistic extreme regions than more sophisticated approaches based on the estimation of the angular density. This requires more investigation which must include suitable criteria.

Fisher's exact test We examined the independence of signs of (ξ_1, ξ_2) by using Fisher's exact test, which is a simple contingency table test for independence. We take "XLF" in the "before" period to illustrate, and use $u = 0.8$ as the threshold. The number of norms above u is 52, and the count (proportion) of the points falling in each of the 4 quadrants (+, +), (+, -), (-, -), (-, +) are summarized in the following contingency table:

Frequency (Percentage)		ξ_2 +	ξ_2 –	Total
ξ_1	+	15 (0.29)	8 (0.15)	23 (0.44)
	–	14 (0.27)	15 (0.29)	29 (0.56)
Total		29 (0.56)	23 (0.44)	52 (1)

Using `R` function `fisher.test`, we get the p value of 0.27. The test results for all the assets are summarized in Tables D.4 and D.5. For all them, no evidence was found that there is a preference for any of the four quadrants. Also, no apparent pattern shows up in the 2×2 tables of probabilities as they change over the four periods "before", "during", "after1" and "after2". This leads us to conclude that the crisis has not caused any shift in the "direction" of the extreme CIDR curves. For example, the extreme scores are roughly equally likely to appear in the (+, +) region before, during and after the crisis.

A more subtle question is if the extreme dependence between the first and second scores has been affected by the crisis. Such an investigation requires rather sophisticated EVT tools, and is the subject of the next section.

3.4 Extremal dependence between the magnitude of the scores

By construction, the scores ξ_j in (2.1) are uncorrelated. Using the methods described below, we determined that there is no extremal dependence between the scores at the

various levels j. There is however extremal dependence between the squares of level 1 and level 2 scores which occurs for certain assets and certain periods. Roughly speaking, such a dependence means that extremely high monotonic trend increases the chance of a pronounced inflection point, i.e. of a large change in the rate of growth or decline. Independence in this sense means that the occurrence of a pronounced inflection point is not associated with a strong upward or downward trend. This section explores this issue in some detail using measures of extremal dependence defined by Coles *et al.* (1999).

Two measures of extremal dependence Suppose two variables X and Y have a common distribution F. We define

$$\chi = \lim_{z \to z_+} P(Y > z \mid X > z),$$

where z_+ is the end-point of F, so that $\chi \in [0, 1]$ is a limiting measure of the tendency for one variable to be large conditional on the other one being large. If $\chi = 0$ the variables X and Y are said to be *asymptotically independent*. If $\chi > 0$, they are said to be *asymptotically dependent*.

More generally, suppose that F_X and F_Y are the marginal distribution functions of X and Y, respectively. We then define

$$\begin{aligned} \chi(q) &= P(F_Y(Y) > q \mid F_X(X) > q) \\ &= P(Y > \mathrm{VaR}_{1-q}(Y) \mid X > \mathrm{VaR}_{1-q}(X)) \end{aligned}$$

and

$$\chi = \lim_{q \to 1} \chi(q).$$

The expression for $\chi(q)$ in terms of VaR shows that χ can be interpreted as the probability that Y exceeds its VaR_α value, given that X exceeds its VaR_α value for extremely small α. Thus, if $\chi = 0.5$, for example, 50% of the time the losses of X are due to extreme losses in Y.

Clearly, χ is a useful measure of extremal dependence within the class of asymptotically dependent distributions, where the value of χ increases with the strength of dependence at extreme levels. A complementary measure $\overline{\chi} \in [-1, 1]$ is defined as follows. For $0 < q < 1$, define

$$\overline{\chi}(q) = \frac{2 \log P(F_X(X) > q)}{\log P(F_X(X) > q, F_Y(Y) > q)} - 1$$

and

$$\overline{\chi} = \lim_{q \to 1} \overline{\chi}(q).$$

When $\overline{\chi} = 1$, the variables X and Y are asymptotically dependent, and when $-1 \leq \overline{\chi} < 1$ they are asymptotically independent. The two measures, χ and $\overline{\chi}$ can thus be used to tests different null hypotheses. Their properties as summarized in the following table, in which independence or dependence refer to these quantities pertaining to extreme events.

Independent	Dependent
$\chi = 0$	$0 < \chi \leq 1$
$-1 \leq \overline{\chi} < 1$	$\overline{\chi} = 1$

Illustrative examples of data models are provided in Appendix C. Generally, extremal independence is more common than extremal dependence. For this reason, one typically first tests the null hypothesis $\overline{\chi} = 1$, and if it is not rejected, assesses the strength of extremal dependence by the value of $\chi > 0$. The details of this procedure are presented in Appendix C.

Application to sector ETFs To illustrate the methodology introduced above, Figure D.11 shows the plots of $\chi(q)$ and $\overline{\chi}(q)$, together with approximate 95% confidence intervals, for the squared scores ξ_1^2 and ξ_2^2 of XLF. Interpretation of these plots is not completely clear because of the large variance of the estimators. However, we can roughly tell that for the "before" "after1" and "after2" periods $\overline{\chi}(q) \to 1$ as $q \to 1$ while for the "during" period this is not the case. In the respective periods, $\chi(q)$ converges to values of around 0.3, 0, 0.3 and 0.4. Furthermore, the likelihood ratio test of asymptotic dependence, see Appendix C, lends support to the use of asymptotically dependent models above a sufficiently high threshold for the "before" "after1" and "after2" periods. Finally, under properly chosen thresholds (we use 80% marginal quantiles as the thresholds for each of the "before" "after1" and "after2" periods), by applying the natural non-parametric estimator (C.5) under the constraint $\hat{\overline{\chi}} = 1$, we obtain $\hat{\chi} = 0.40$ for the "before" period, $\hat{\chi} = 0.43$ for the "after1" period and $\hat{\chi} = 0.35$ for the "after2" period. After accounting for the variance, these results are consistent with the apparent stable levels shown in Figure D.11.

The estimation results for all the nine ETFs are shown in Table D.6. For each ETF, the columns "q" and "n_u", respectively, give the quantile of the threshold and the number of observations exceeding that threshold during each period. We also give an estimate of $\overline{\chi}$ with corresponding standard errors. The "Pvalue" entries are the results of the likelihood ratio test of asymptotic dependence, with the null hypothesis $\overline{\chi} = 1$. Only if the results are insignificant, do we give a non-zero estimate of χ. The estimation results are sensitive to the choice of the threshold. We choose it carefully in a multi-stage process. First, we look at the the mean residual life plot using the `R` function `mrlplot`. Next, we use the `R` function `tcplot`, which is a plot of parameter estimates at various thresholds for PoT modeling. The valid threshold will be a value such that the parameter estimates depicted are approximately constant above the threshold. We tried several thresholds suggested by the above two approaches and selected a proper quantile q such that the estimation results become stable and reasonable. (An inappropriately chosen q, hence the threshold u, can result in very unstable estimation.)

After accounting for sampling errors, we can clearly observe a pattern among most of the ETF sectors (XLK, XLY, XLV, XLI and XLB) that the "before" periods show a

strong signal of extremal dependence between ξ_1^2 and ξ_2^2 while for some other ETF sectors (XLE and XLU) it occurs in the "after2" period. Since the "after2" is similar to the "before" period in other respects, we can conclude that during the crisis as well as the short period right after the crisis, there tends to be no asymptotic dependence between ξ_1^2 and ξ_2^2 for any ETF sectors. Generally, the extremal dependence is of a different type "before" than in the other three periods, or is of a different type in the "during" period.

3.5 Summary and concluding Remarks

We employed the extreme value theory *peak-over-threshold* (PoT) method and its extensions to study the extremal behavior of the cumulative intraday returns (CIDRs) of a set of financial assets. The CIDRs are curves which describe how the return evolves over a trading day. We defined and estimated risk measures such as Value-at-Risk (return level) and expected shortfall for such function valued data. The norm of the CIDRs can be viewed as a new measure of intraday price movements which incorporates intraday variability and systematic trends or changes in direction (e.g. from growth to decline). The scores quantify the presence of specific patterns, like trends or reversals. We proposed several ways of visualizing the risk associated with CIDR's.

Our major findings can be summarized as follows:

1. Scores of the first three FPCs of a CIDR series typically contribute $90\% \sim 95\%$ of the total variability. The first score, ξ_1, typically contributes more than 85%. This indicates that ξ_1 could be a more informative summary of the daily returns than the conventional point-to-point returns based on the closing price. (For example, large ξ_1 indicates that the price increases almost monotonically throughout the day.)

2. The norm of the CIDRs, $\|R\|$, has an exponential-decaying tail when the estimation is performed over a relatively short period (1-2 years). This is likely because the clustering phenomenon is not pronounced over such time periods. The intraday cumulative returns and their variability are either large or small, so the tail, which measures the presence of observations much larger than most observations, is light. This is in contrast to point–to–point returns over long periods of time, which are often modeled as heavy-tailed functions.

3. Estimation of the risk measures leads to the following general pattern around the financial crisis of 2008:

$$\text{during} \;\geq\; \text{after1} \;\geq\; \text{after2} \;\approx\; \text{before},$$

   where the periods are arranged from the most risky to the least risky. Since each period has the length of about a year and a half, this implies the crisis had a monotonically decaying, long-term impact on the intraday market risk extending

beyond the commonly accepted end of the crisis. The pattern of the decay of risk measures based on point–to–point returns is often not monotonic.

4. Risk measures of CIDRs, in general, contain different information about the extreme behavior of returns than those measures computed for point-to-point returns. The impact of the crisis of 2008 on extreme risk is different for the CIDRs and point-to-point returns, both in terms of a general pattern and impact on specific sectors.

5. The crisis impacted extremal dependence between the volatilities of the principal components of the shapes of the CIDRs.

There is clearly room for improvement and extensions of the methodology we proposed. Multivariate EVT could in principle be applied to model the joint extremal behavior of the scores, some effort in this direction is outlined in Appendix B, but more sophisticated and accurate approaches might be possible. Rather than using the PoT approach, likelihood based methods that incorporate temporal dependence might be feasible. Potential conditional heteroskedasticity might possibly be incorporated; some research in this direction has been done, Aue *et al.* (2016), Kokoszka *et al.* (2017b), Cerovecki *et al.* (2018), but it does not offer any guidance on how to incorporate these models into the study of extremal behavior of CIDR curves. In all, we hope that our exploration will be received with some interest as an attempt to quantify extreme behavior of samples of curves with a potential to quantify some aspects of investment risk.

References

Aue, A., Horváth, L. and Pellatt, D. (2016). Functional generalized autoregressive conditional heteroskedasticity. *Journal of Time Series Analysis*, **38,** 3–21.

Bensalah, Y. (2000). Steps in applying extreme value theory to finance: A review. *Bank of Canada Working Paper*, **101,** 2000–20.

Black, F. (1972). Capital market equilibrium with restricted borrowing. *Journal of Business*, **45,** 444–454.

Bosq, D. (2000). *Linear Processes in Function Spaces*. Springer.

Brodin, E. and Klüppelberg, C. (2008). Extreme value theory in finance. *Encyclopedia of Quantitative Risk Analysis and Assessment*, **2**.

Brooks, C., Clare, A. D., Molle, J. W. Dalle and Persand, G. (2005). A comparison of extreme value theory approaches for determining value-at-risk. *Journal of Empirical Finance*, **12,** 339–352.

Cerovecki, C., Francq, C., Hörmann, S. and Zakoian, J-M. (2018). Functional GARCH models: the quasi-likelihood approach and its applications. Technical Report. Université libre de Bruxelles.

Coles, S. G., Heffernan, J. and Tawn, J. A. (1999). Dependence measures for extreme value analysis. *Extremes*, **3,** 5–38.

Danielsson, J. and de Vries, C. G. (2000). Value-at-risk and extreme returns. *Annales d'Economie et de Statistique*, **60,** 239–270.

Davison, A. C. and Smith, R. L. (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodological)*, **52,** 393–442.

Dionne, G., Duchesne, P. and Pacurar, M. (2009). Intraday value at risk (ivar) using tick-by-tick data with application to the Toronto Stock Exchange. *Journal of Empirical Finance*, **16,** 777–792.

Embrechts, P., Klüppelberg, C. and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer, Berlin.

Gençay, R. and Selçuk, F. (2004). Extreme value theory and value-at-risk: Relative performance in emerging markets. *International Journal of Forecasting*, **20,** 287–303.

Gençay, R., Selçuk, F. and Ulugülyağci, A. (2003). High volatility, thick tails and extreme value theory in value-at-risk estimation. *Insurance: Mathematics and Economics*, **33,** 337–356.

Gilli, M. and Këllezi, E. (2006). An application of extreme value theory for measuring financial risk. *Computational Economics*, **27,** 207–228.

Giot, P. (2005). Market risk models for intraday data. *European Journal of Finance*, **11,** 309–324.

Hall, P., Watson, G. S. and Cabrera, J. (1987). Kernel density estimation with spherical data. *Biometrika*, **74,** number 4, 751–762.

Heffernan, J. (2000). A directory of coefficients of tail dependence. *Extremes*, **3:3,** 279–290.

Horváth, L. and Kokoszka, P. (2012). *Inference for Functional Data with Applications*. Springer.

Horváth, L., Kokoszka, P. and Rice, G. (2014). Testing stationarity of functional time series. *Journal of Econometrics*, **179,** 6682.

Kokoszka, P., Miao, H. and Zhang, X. (2015). Functional dynamic factor model for intraday price curves. *Journal of Financial Econometrics*, **13,** 456–477.

Kokoszka, P., Miao, H. and Zheng, B. (2017). Testing for asymmetry in betas of cumulative returns: Impact of the financial crisis and crude oil price. *Statistics & Risk Modeling*, **34,** 33–53.

Kokoszka, P. and Reimherr, M. (2013). Predictability of shapes of intraday price curves. *The Econometrics Journal*, **16,** 285–308.

Kokoszka, P. and Reimherr, M. (2017). *Introduction to Functional Data Analysis*. CRC Press.

Kokoszka, P., Rice, G. and Shang, H. L. (2017b). Inference for the autocovariance of a functional time series under conditional heteroscedasticity. *Journal of Multivariate Analysis*, **162,** 35–50.

Kourouma, L., Dupre, D., Sanfilippo, G. and Taramasco, O. (2011). Extreme value at risk and expected shortfall during financial crisis. Technical Report. HAL.

Ledford, A. and Tawn, A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, **83,** 169–187.

Ledford, A. and Tawn, A. (1998). Concomitant tail behavior for extremes. *Advances in Applied Probability*, **30,** 197–215.

Lintner, J. (1965). The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, **47,** 13–37.

Lucca, D. O. and Moench, E. (2015). The pre-FOMC announcement drift. *The Journal of Finance*, **70,** 329–371.

Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments*. John Wiley.

McNeil, A., Frey, R. and Embrechts, P. (2005). *Quantatitative Risk Management*. Princeton University Press.

McNeil, A. J. and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach. *Journal of Empirical Finance*, **7,** 271–300.

Neftci, S. N. (2000). Value-at-risk calculations, extreme events, and tail estimation. *Journal of Derivatives*, **7(3),** 23–38.

Poon, S., Rockinger, M. and Tawn, J. (2004). Extreme value dependence in financial markets: Diagnostics, models, and financial implications. *The Review of Financial Studies*, **17,** 581–610.

Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer.

Resnick, S. I. (2007). *Heavy–Tail Phenomena*. Springer, New York.

Rocco, M. (2014). Extreme value theory for finance: a survey. *Journal of Economics Surveys*, **28,** 82–108.

Sharpe, W. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance*, **19,** 425–442.

Stephenson, A. (2012). Statistics of multivariate extremes. Technical Report. CRAN.

Zhang, X. (2016). White noise testing and model diagnostic checking for functional time series. *Journal of Econometrics*, **194,** number 1, 76–95.

Appendices

A Variances of $\widehat{\mathrm{VaR}}_\alpha$ and $\widehat{\mathrm{ES}}_\alpha$

Denote $S_u = P(Y > u)$ and $\hat{S}_u = N_u/N$. Observe that $\hat{S}_u$ is the maximum likelihood estimate of S_u, as the number of exceedances of u follows the Binomial distribution $\sim \mathrm{Bin}(n, S_u)$. Since both $\widehat{\mathrm{VaR}}_\alpha$ and $\widehat{\mathrm{ES}}_\alpha$ can be viewed as a function of $\hat{\sigma}$, $\hat{\gamma}$ and $\hat{S}_u$, i.e.

$$\widehat{\mathrm{VaR}}_\alpha := g_1(\hat{S}_u, \hat{\sigma}, \hat{\gamma}) = u + \frac{\hat{\sigma}}{\hat{\gamma}}\left[\left(\frac{\hat{S}_u}{\alpha}\right)^{\hat{\gamma}} - 1\right]$$

and

$$\widehat{\mathrm{ES}}_\alpha := g_2(\hat{S}_u, \hat{\sigma}, \hat{\gamma}) = u + \frac{\hat{\sigma}}{1-\hat{\gamma}} + \frac{\hat{\sigma}}{\hat{\gamma}(1-\hat{\gamma})}\left[\left(\frac{\hat{S}_u}{\alpha}\right)^{\hat{\gamma}} - 1\right],$$

large sample standard errors or confidence intervals for VaR_α and ES_α can be derived from the delta method. From standard properties of the binomial distribution, $\mathrm{VaR}(\hat{S}_u) \approx \hat{S}_u(1-\hat{S}_u)/n$ and denote the (i,j) term of the variance-covariance matrix of $\hat{\sigma}$ and $\hat{\gamma}$ by $v_{i,j}$, then the complete covariance-variance matrix for $(\hat{\sigma}, \hat{\gamma}, \hat{S}_u)$ is approximately

$$V = \begin{bmatrix} \hat{S}_u(1-\hat{S}_u)/n & 0 & 0 \\ 0 & v_{1,1} & v_{1,2} \\ 0 & v_{2,1} & v_{2,2} \end{bmatrix}$$

Hence, by the delta method,

$$\mathrm{Var}(\widehat{\mathrm{VaR}}_\alpha) \approx \bigtriangledown g_1^T V \bigtriangledown g_1$$

and

$$\mathrm{Var}(\widehat{\mathrm{ES}}_\alpha) \approx \bigtriangledown g_2^T V \bigtriangledown g_2,$$

where

$$\begin{aligned}\bigtriangledown g_1^T &= \left[\frac{\partial g_1}{\partial S_u},\ \frac{\partial g_1}{\partial \sigma},\ \frac{\partial g_1}{\partial \gamma}\right] \\ &= \left\{\frac{\sigma S_u^{\gamma-1}}{\alpha^\gamma},\ \frac{1}{\gamma}\left[\left(\frac{S_u}{\alpha}\right)^\gamma - 1\right],\ -\frac{\sigma}{\gamma^2}\left[\left(\frac{S_u}{\alpha}\right)^\gamma - 1\right] + \frac{\sigma}{\gamma}\left(\frac{S_u}{\alpha}\right)^\gamma \log\left(\frac{S_u}{\alpha}\right)\right\}\end{aligned}$$

and

$$\begin{aligned}\bigtriangledown g_2^T &= \left[\frac{\partial g_2}{\partial S_u},\ \frac{\partial g_2}{\partial \sigma},\ \frac{\partial g_2}{\partial \gamma}\right] \\ &= \left\{\frac{\sigma S_u^{\gamma-1}}{(1-\gamma)\alpha^\gamma},\ \frac{1}{\gamma(1-\gamma)}\left[\gamma + \left(\frac{S_u}{\alpha}\right)^\gamma - 1\right],\right. \\ &\qquad \left.\frac{\sigma}{(1-\gamma)^2} + \frac{\sigma(2\gamma-1)}{\gamma^2(1-\gamma)^2}\left[\left(\frac{S_u}{\alpha}\right)^\gamma - 1\right] + \frac{\sigma}{\gamma(1-\gamma)}\left(\frac{S_u}{\alpha}\right)^\gamma \log\left(\frac{S_u}{\sigma}\right)\right\}\end{aligned}$$

evaluated at $(\hat{S}_u, \hat{\sigma}, \hat{\gamma})$.

B Estimation of the angular density

This section explains how the density of $\mathbf{Y}(u)$ in (3.5) can be estimated nonparametrically. Since we consider two scores, this is a density on $(-\pi, \pi)$. The idea is that if $\theta_i, i = 1, \ldots, n$, are *i.i.d.* points from a density f, then its kernel estimator is given by

$$\hat{f}_h(\theta) := \frac{2}{n}\sum_{i=1}^{n} \phi_h(\tan(\theta - \theta_i)/2))/\cos^2((\theta - \theta_i)/2), \tag{B.1}$$

where $\phi_h(x) = \phi(x/h)/h$ is probability density on the real line. The idea behind this estimator is to lift a distribution on the real line to the unit circle via the transformation $\theta \mapsto \tan(\theta/2)$. By default, we will use the standard normal density for ϕ (see also (Hall *et al.*, 1987)). To apply this idea to the procedure of resampling angles, suppose $(\xi_{1,i}, \xi_{2,i}), i = 1, \ldots, N_u$ are the points with norm greater than the threshold u and define $\theta_i := \arctan(\xi_{2,i}/\xi_{1,i})$ for $i = 1, \ldots, N_u$. Choosing the bandwidth $h = 2\pi/N_u^{2/3}$, we compute (B.1) using the $\theta_i's$, and so obtain an estimated density function $\hat{f}_h(\theta)$ of the angles θ on $(-\pi, \pi)$. From this density, we sample angles that can be used to generate new extreme points. Figure B.1 illustrates this method using the asset XLF in the "Before" period.

C Elaboration on the measures χ and $\overline{\chi}$

This section provides details of inference based on the extreme dependence measures χ and $\overline{\chi}$. Essentially, χ provides a measure with which to describe the strength of dependence

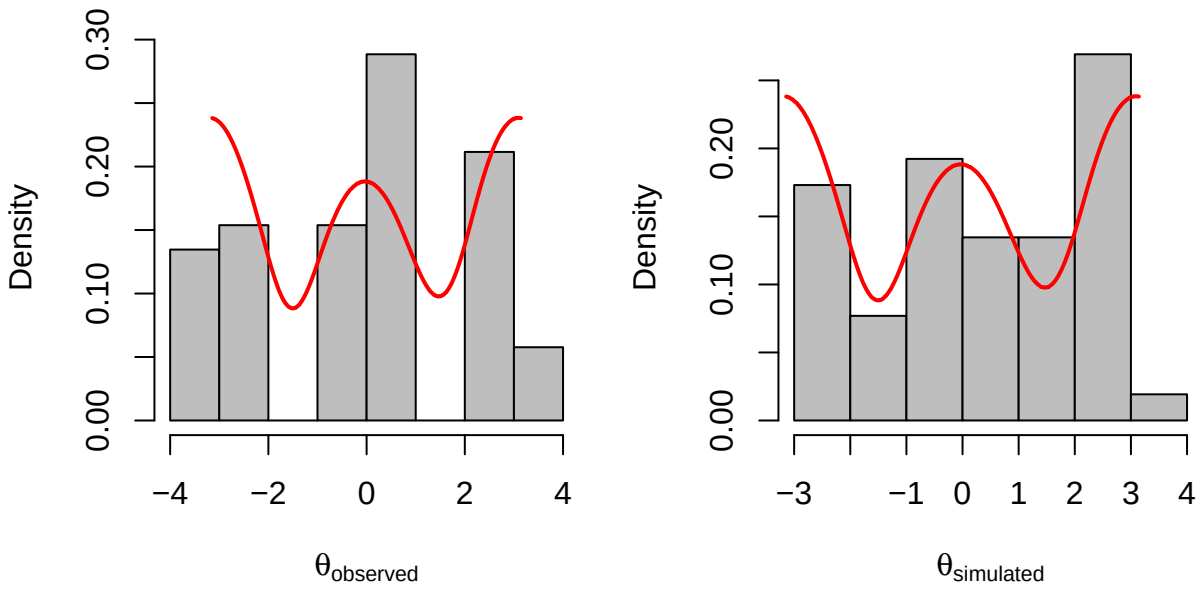

Figure B.1: Left panel: histogram of the observed $\theta_i's$ with the estimate $\hat{f}_h(\theta)$ superposed. Right panel: histogram of the same number of new $\theta_i's$ generated from $\hat{f}_h(\theta)$.

within the class of asymptotically dependent variables, while $\overline{\chi}$ provides a corresponding measure within the class of asymptotically independent variables. Taken together, the pair $(\chi, \bar{\chi})$ provides more complete information characterizing the form and the degree of extremal dependence of any two random variables. For asymptotically dependent variables, $\bar{\chi} = 1$ and the value of $\chi > 0$ measures the strength of dependence. For asymptotically independent variables, $\chi = 0$, and the value of $\overline{\chi} \in [-1, 1)$ might be used to quantify the strength of extremal independence, which may be interpreted also as a secondary or hidden tail dependence in terms of the notion of *hidden regular variation* (Resnick, 2007).

We begin by providing two examples where $\chi > 0$ (asymptotic dependence) and $\chi = 0$ (asymptotic independence).

Example C.1 In this example, X and Y have a common risk factor. Let Z, ε_x and ε_y be independent GPD$(\gamma, \gamma, 0)$ random variables with the same index $\gamma > 0$, such that

$$\mathbb{P}(Z > u) = \mathbb{P}(\varepsilon_x > u) = \mathbb{P}(\varepsilon_y > u) = (1 + u)^{-1/\gamma},$$

for $u > 0$. Consider the simple factor model

$$X = \sigma_z Z + a\varepsilon_x \qquad \text{and} \qquad Y = \sigma_z Z + b\varepsilon_y,$$

where σ_z, a and b are positive. By applying Lemma C.1 below with $a_1 = b_1 = \sigma_z$, $a_2 = a$,

$b_2 = 0$, and $a_3 = 0$, $b_3 = b$, we obtain

$$\sigma_X^{1/\gamma} = \sigma_z^{1/\gamma} + a^{1/\gamma}, \qquad \sigma_Y^{1/\gamma} = \sigma_z^{1/\gamma} + b^{1/\gamma},$$

and

$$\chi(X,Y) = \frac{\sigma_z^{1/\gamma}}{\sigma_z^{1/\gamma} + a^{1/\gamma}} \wedge \frac{\sigma_z^{1/\gamma}}{\sigma_z^{1/\gamma} + b^{1/\gamma}} = \frac{1}{1 + \max(a^{1/\gamma}, b^{1/\gamma})/\sigma_z^{1/\gamma}}.$$

Thus, the larger the ratio $\max(a^{1/\gamma}, b^{1/\gamma})/\sigma_z^{1/\gamma}$, the smaller the extremal dependence. This is natural since in this case, σ_z is relatively small and the common factor $\sigma_z Z$ has relatively smaller contribution to the simultaneous extremes of X and Y. On the other hand, if σ_z dominates a and b, the common risk factor leads to more frequent simultaneous extremes, which is reflected by relatively larger values of the extremal dependence measure χ.

EXAMPLE C.2 Suppose now that (X, Y) are jointly Gaussian with correlation coefficient $\rho \in (-1, 1)$. It can be shown that $\chi = 0$ in this case, i.e. X and Y are always asymptotically independent (unless $\rho = 1$). The measure $\bar{\chi}$ is more appropriate in this case and $\overline{\chi} = \rho$. See Heffernan (2000) for the details.

LEMMA C.1 *Let $Z_1, \ldots, Z_p$ be independent heavy–tailed GPD$(\gamma, \gamma, 0)$ random variables, i.e., $\mathbb{P}(Z_i > u) = (1+u)^{-1/\gamma}$, $u > 0$, for some $\gamma > 0$. Define*

$$X = \sum_{i=1}^{p} a_i Z_i \quad \text{and} \quad Y = \sum_{i=1}^{p} b_i Z_i,$$

where $a_i > 0$ and $b_j > 0$, for some $i, j \in \{1, \ldots, p\}$. Then,

$$\mathbb{P}(X > u,\ Y > u) \sim u^{-1/\gamma} \sum_{i=1}^{p} (a_i \wedge b_i)_+^{1/\gamma}, \quad \text{as } u \to \infty, \tag{C.1}$$

where $x_n \sim y_n$ means $x_n/y_n \to 1$, $x \wedge y := \min\{x, y\}$ and $(x)_+ := \max\{x, 0\}$.

Consequently, the tail dependence coefficient between X and Y equals

$$\chi(X,Y) = \sum_{i=1}^{p} \left(\frac{a_i}{\sigma_X} \wedge \frac{b_i}{\sigma_Y} \right)_+^{1/\gamma}, \tag{C.2}$$

where $\sigma_X^{1/\gamma} = \sum_{i=1}^{p} (a_i)_+^{1/\gamma}$ and $\sigma_Y^{1/\gamma} = \sum_{i=1}^{p} (b_i)_+^{1/\gamma}$.

PROOF: We shall first obtain Relation (C.2) as a consequence of (C.1). For simplicity, let $\alpha := 1/\gamma > 0$. By formally applying (C.1) to $a_i := b_i$, we obtain

$$\mathbb{P}(Y > u) \sim \sigma_Y^{\alpha} u^{-\alpha}, \quad \text{where } \sigma_X^{\alpha} = \sum_{i=1}^{p} (b_i)_+^{\alpha},$$

and similarly $\mathbb{P}(X > u) \sim \sigma_X^\alpha u^{-\alpha}$, as $u \to \infty$. Observe that $\sigma_X > 0$ and $\sigma_Y > 0$ since $\max_{i=1,\dots,p} a_i > 0$ and $\max_{i=1,\dots,p} b_i > 0$.

Thus, $\mathbb{P}(X/\sigma_X > u) \sim \mathbb{P}(Y/\sigma_Y > u) \sim u^{-\alpha}$, as $u \to \infty$, and hence

$$\chi(X,Y) = \lim_{u\to\infty} \mathbb{P}(X/\sigma_X > u | Y/\sigma_Y > u) = \lim_{u\to\infty} u^\alpha \mathbb{P}(X > \sigma_X u, Y > \sigma_Y u).$$

Now, by applying (C.1) with a_i replaced by a_i/σ_X and b_i replaced by b_i/σ_Y, the formula (C.2) follows.

Relation (C.1) is well known but we give a proof for completeness. This result can be understood in terms of the so-called *one big jump* heuristic. Namely, since the Z_i's are independent and heavy–tailed with the same tail exponent $\alpha = 1/\gamma > 0$, asymptotically, a linear combination of the Z_i's is large if one and only one of its terms is extreme. That is, the probabilities that two or more components $a_i Z_i$ contribute to an extreme value of X is asymptotically negligible. Using this principle, one can intuitively see that

$$\mathbb{P}(X > u, Y > u) \sim \sum_{i=1}^p \mathbb{P}(a_i Z_i > u, b_i Z_i > u) = u^{-\alpha} \sum_{i=1}^p (a_i \wedge b_i)_+^\alpha,$$

as $u \to \infty$.

We will make the above heuristic precise using the notion of multivariate regular variation. We start with some terminology. A set $A \subset \mathbb{R}^p \setminus \{\mathbf{0}\}$ is said to be bounded away from $\mathbf{0}$ if there us a ball $B(\mathbf{0}, \epsilon)$ centered at the origin with radius $\epsilon > 0$ such that $A \cap B(\mathbf{0}, \epsilon) = \emptyset$. A random vector $\boldsymbol{Z} = (Z_1, \dots, Z_p)^\top$ is said to be *regularly varying* in $\mathbb{R}^p \setminus \{\mathbf{0}\}$ if there exists a Borel measure ν on $\mathbb{R}^p \setminus \{\mathbf{0}\}$, such that

$$c(u)\mathbb{P}(\boldsymbol{Z} \in uA) \to \nu(A), \tag{C.3}$$

for all measurable bounded away from $\mathbf{0}$ and ν-continuity sets A, i.e., such that $\nu(\partial A) = 0$, where $\partial A = \overline{A} \setminus A^\circ$ is the boundary of A.

In the simple case above, it is easy to see that the vector $\boldsymbol{Z}$ with independent GPD$(\gamma, \gamma, 0)$ components is regularly varying, where ν is supported on the positive orthant $[0,\infty)^p$. One can identify the measure ν by taking $A = [\mathbf{0}, \boldsymbol{x}]^c = \{\boldsymbol{y} \in \mathbb{R}^p \ : \ y_i > x_i \text{ for some } i = 1, \dots, p\}$, for $\boldsymbol{x} \geq \mathbf{0}$. Indeed, by using the independence of the Z_i's and the inclusion-exclusion formula, we get as $u \to \infty$,

$$\mathbb{P}(\boldsymbol{Z} \in uA) \sim \mathbb{P}\left(\cup_{i=1}^p \{Z_i > u x_i\}\right) \sim u^{-\alpha} \sum_{i=1}^p x_i^\alpha =: u^{-\alpha}\nu(A).$$

Thus (C.3) holds with $c(u) := u^\alpha$, where the measure ν is concentrated on the positive axes $\ell_i := \{\lambda \boldsymbol{e}_i \ : \ \lambda > 0\}$, $i = 1, \dots, p$, with $\boldsymbol{e}_1 = (1, 0, \dots, 0)^\top, \dots, \boldsymbol{e}_p = (0, \dots, 0, 1)^\top$ denoting the standard basis of $\mathbb{R}^p$. In general, for a measurable set A, we have the formula

$$\nu(A) = \sum_{i=1}^p \nu_\alpha(\pi_i(A \cap \ell_i)), \tag{C.4}$$

where $\nu_\alpha(x,\infty) = x^{-\alpha},\ x > 0$ is a measure on the positive half-line $(0,\infty)$, and $\pi_i : \mathbb{R}^p \to \mathbb{R}$ is the projection on the i-th coordinate axis.

With this general tool, we can establish the tail behavior of various functionals of $\boldsymbol{Z}$ by relating them to suitable sets A. In particular, observe that

$$\{X > u\} = \{\boldsymbol{Z} \in uA\}, \quad \text{and} \quad \{Y > u\} = \{\boldsymbol{Z} \in uB\}.$$

where $A := \{\boldsymbol{z} = (z_i)_{i=1}^p : \sum_{i=1}^p a_i z_i > 1\}$ and $B := \{\boldsymbol{z} : \sum_{i=1}^p b_i z_i > 1\}$. Observe that both A and B are bounded away from $\mathbf{0}$ and it can be shown, using the scaling properties of the measure ν, that A and B are both ν-continuity sets. Thus, by (C.3) applied to the set $A \cap B$, we obtain

$$u^\alpha \mathbb{P}(X > u, Y > u) = u^\alpha \mathbb{P}(\boldsymbol{Z} \in u(A \cap B)) \to \nu(A \cap B),$$

as $u \to \infty$. By elementary geometric considerations, however, $(A \cap B) \cap \ell_i = \emptyset$, unless both a_i and b_i are positive. In this case, if $a_i \wedge b_i > 0$, we have

$$\pi_i(A \cap B \cap \ell_i) = (1/a_i, \infty) \cap (1/b_i, \infty) = (1/(a_i \wedge b_i), \infty).$$

This, in view of (C.4), yields the formula (C.1) and completes the proof. ∎

Equivalent definition of χ Before we discuss any estimation methods of χ and $\overline{\chi}$, it is useful to introduce an alternative way of defining the two measures. The joint distribution of a set of variables can be separated into their respective marginal distributions and dependence structure among them. This idea is also well known in the study of *copulas*. In order to focus on the dependence structure of two variables, it is helpful to remove the influence of marginal aspects first by transforming the raw data to a common marginal distribution. After such a transformation, differences in distributions are purely due to dependence structures. Now we transform the bivariate variables (X, Y) to unit Fréchet marginals S and T as follows:

$$S = -1/\log F_X(X) \qquad \text{and} \qquad T = -1/\log F_Y(Y),$$

where F_X and F_Y are the marginal distribution functions of X and Y respectively. Notice that there are two typical ways to estimate F_X and F_Y practically. Each of the two marginals can either be approximated from a GPD family using univariate extreme value theory, or be simply approximated by its empirical cumulative distribution function (ECDF). It follows that S and T have the common distribution function $F(s) = e^{-1/s}$. Thus, we can show that χ can be equivalently defined as

$$\begin{aligned}
\chi &= \lim_{q \to 1} P(F(T) > q \mid F(S) > q) \\
&= \lim_{s \to \infty} P(T > s \mid S > s) \\
&= \lim_{s \to \infty} \frac{P(T > s, S > s)}{P(S > s)}.
\end{aligned}$$

Estimation of χ and $\overline{\chi}$ Ledford and Tawn (1996, 1998) established that under weak conditions

$$P(S > s, T > s) \sim \mathcal{L}(s)s^{-1/\eta} \quad \text{as } s \to \infty,$$

where $0 < \eta \leq 1$ is a constant and $\mathcal{L}$ is a slowly varying function. From this representation we have

$$\overline{\chi} = 2\eta - 1$$

and if $\overline{\chi} = 1$, corresponding to $\eta = 1$, then $\chi = \lim_{s\to\infty} \mathcal{L}(s)$. So the estimation of η and $\lim_{s\to\infty} \mathcal{L}(s)$ provide the basis for estimating χ and $\overline{\chi}$. Here, the key point is to estimate the joint distribution $P(S > s, T > s)$ hence η and $\mathcal{L}(s)$. Let $Z = \min(S, T)$, so $P(S > s, T > s) = P(\min(S, T) > s) = P(Z > s)$. Inference follows using univariate extreme value techniques to fit a generalized Pareto distribution to the data points in Z that exceed a large fixed threshold u, then the estimated shape parameter of the fitted distribution provides an estimate of η.

Before we estimate χ, it is important to decide if there exists an asymptotic dependence. We thus first test the null hypothesis $\overline{\chi} = 1$. Only if there is no significant evidence to reject it, we estimate χ.

Testing $\overline{\chi} = 1$ There are basically two major different ways to test if $\overline{\chi} = 1$. First, we can simply create a confidence interval for an estimate of $\overline{\chi}$. From the discussion above, the estimate and standard error of the shape parameter η can be obtained in the usual way from standard likelihood theory. So by the Delta method, one can obtain the corresponding estimate and standard error of $\overline{\chi}$ as

$$\hat{\overline{\chi}} = 2\hat{\eta} - 1 \qquad \text{and} \qquad SE(\hat{\overline{\chi}}) = 2SE(\hat{\eta})$$

where $SE(\hat{\eta})$ can be obtained from the standard maximum likelihood method. Hence a 95% confidence interval of the true $\overline{\chi}$ is approximated by

$$\left(\hat{\overline{\chi}} - 1.96SE(\hat{\overline{\chi}}), \quad \hat{\overline{\chi}} + 1.96SE(\hat{\overline{\chi}})\right)$$

as the sample size in the study is big enough. Thus, the hypothesis $\overline{\chi} = 1$ is rejected (at level 0.05) when the confidence interval does not capture 1. For the second way, given that $\overline{\chi} = 1$ corresponds to $\eta = 1$, one can use `anova` in `R` to perform a likelihood ratio test for asymptotic dependence, with the null hypothesis $\eta = 1$ versus the alternative $\eta < 1$. We refer to (Stephenson, 2012) for the details of `R` implementation.

Estimation of χ Once the above test shows no significant evidence to reject $\overline{\chi} = 1$, we can estimate χ under the assumption that $\overline{\chi} = \eta = 1$. (Poon *et al.*, 2004) provided a maximum-likelihood estimator of χ, i.e. the natural non-parametric estimator under the constraint $\hat{\overline{\chi}} = 1$

$$\hat{\chi} = \frac{un_u}{n}, \tag{C.5}$$

$$SE(\hat{\chi}) = \sqrt{\frac{u^2 n_u (n - n_u)}{n^3}}$$

where n is the sample size and n_u is the number of observations of variable Z that exceed the threshold u.

D Large tables and displays

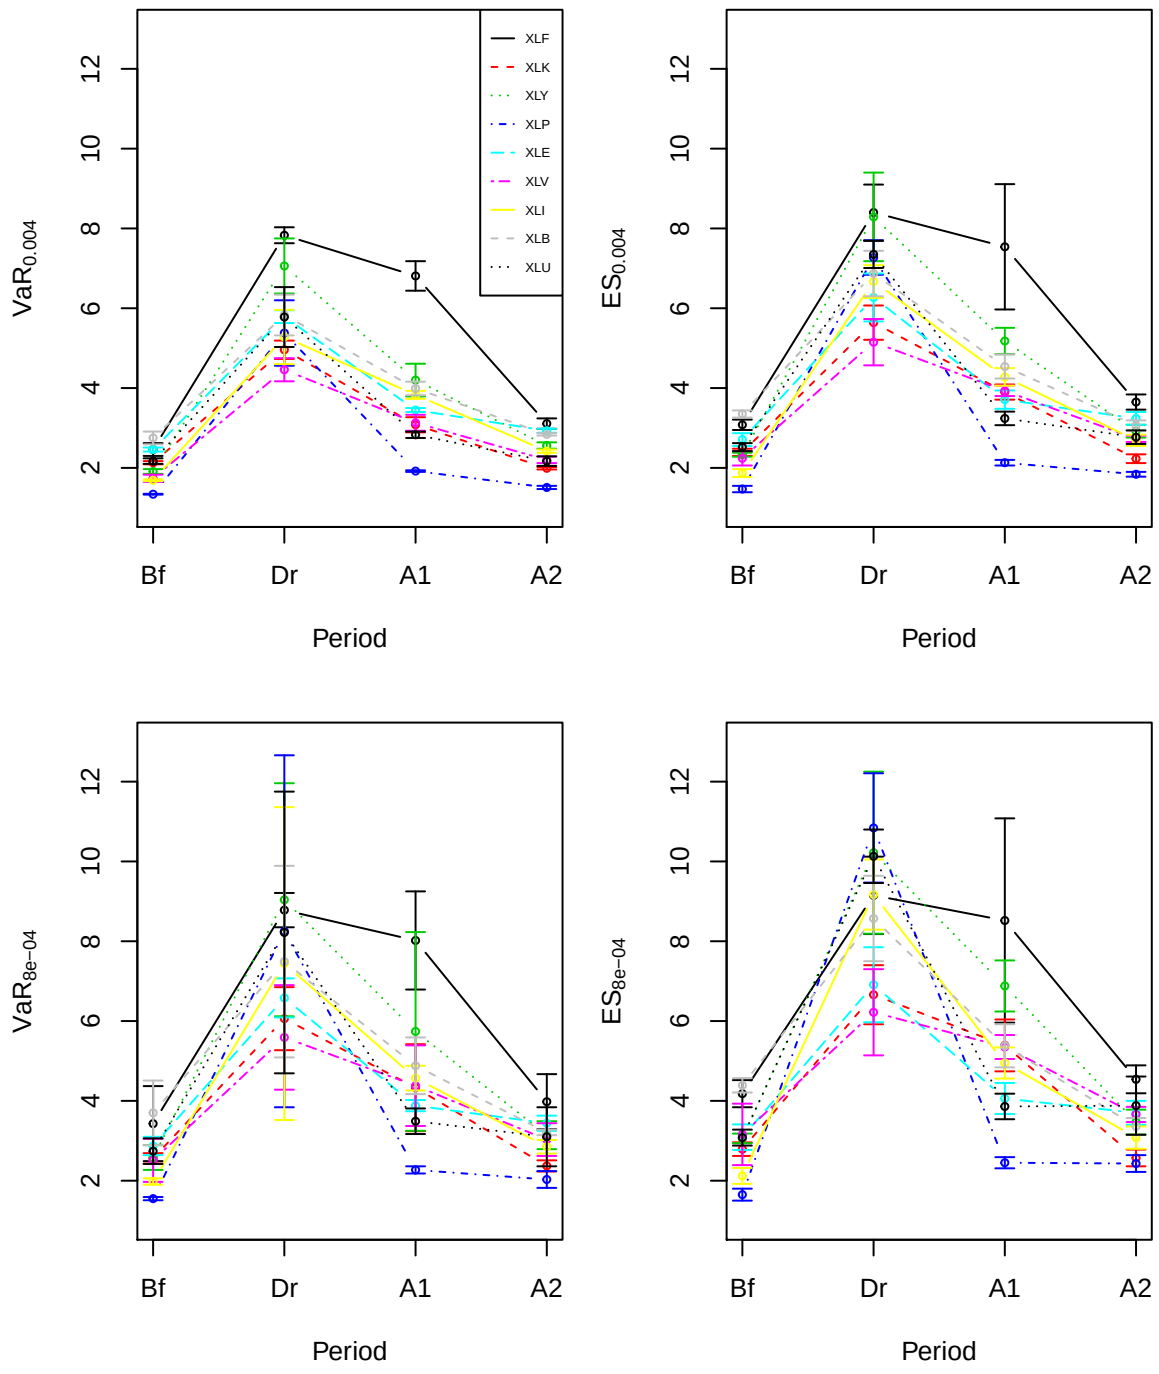

Figure D.1: VaR & ES Estimates (with standard errors) of norms of CIDRs for the nine Sector EFTs.

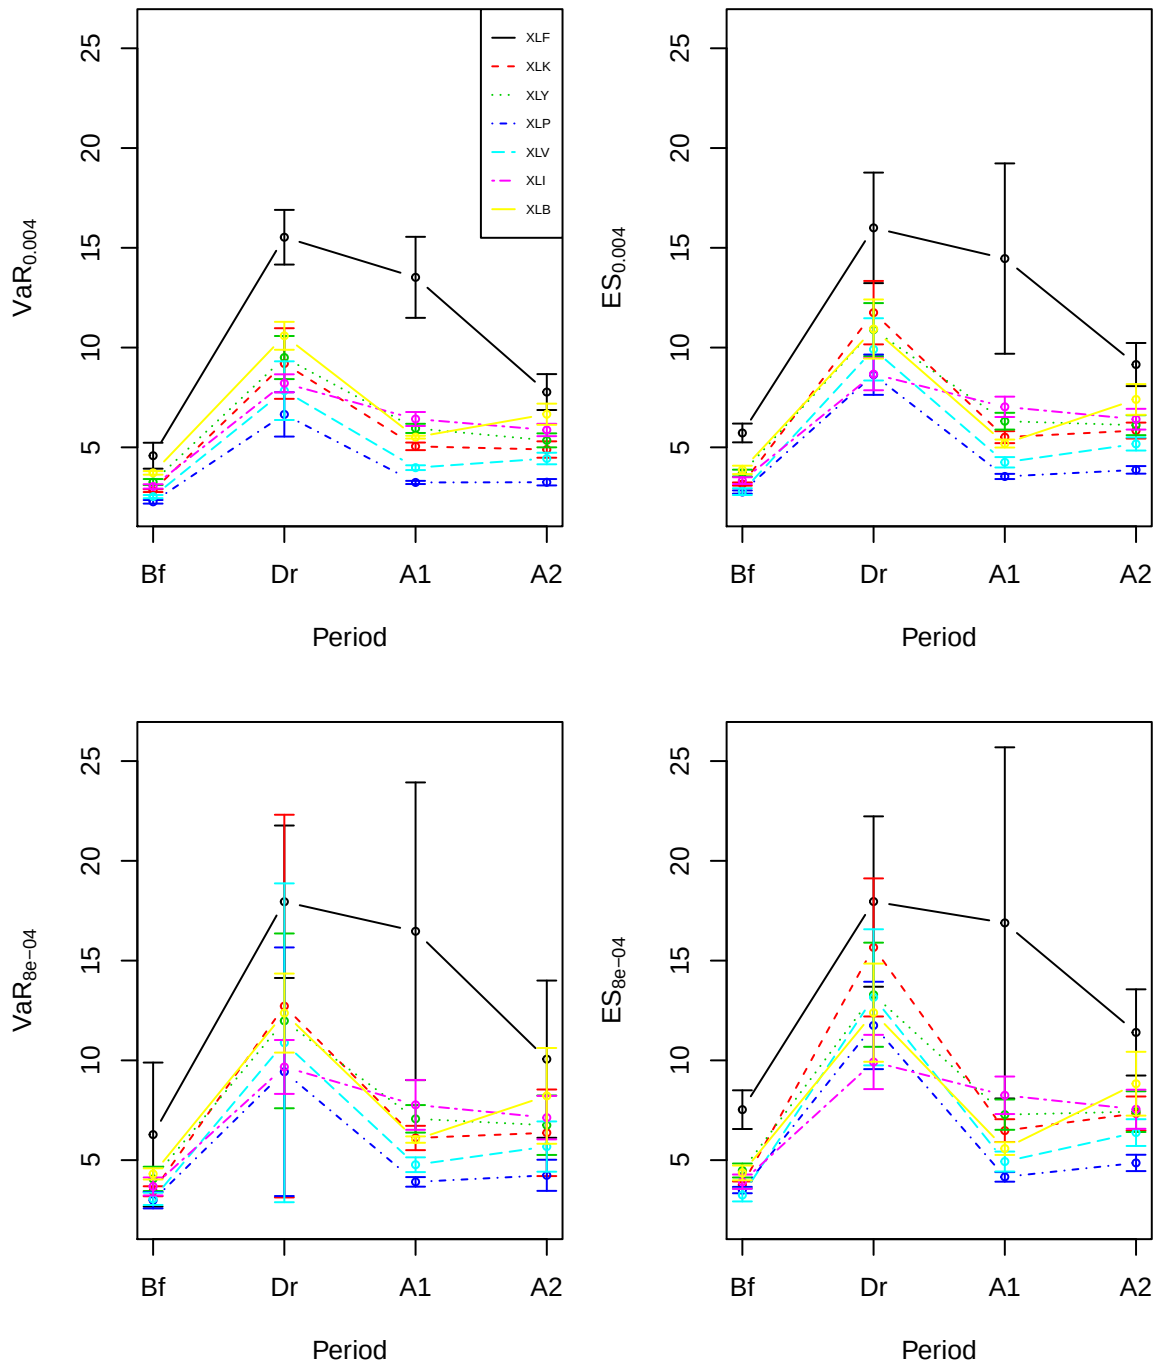

Figure D.2: VaR & ES Estimates (with standard errors) of magnitude of point-to-point returns for the seven Sector EFTs.

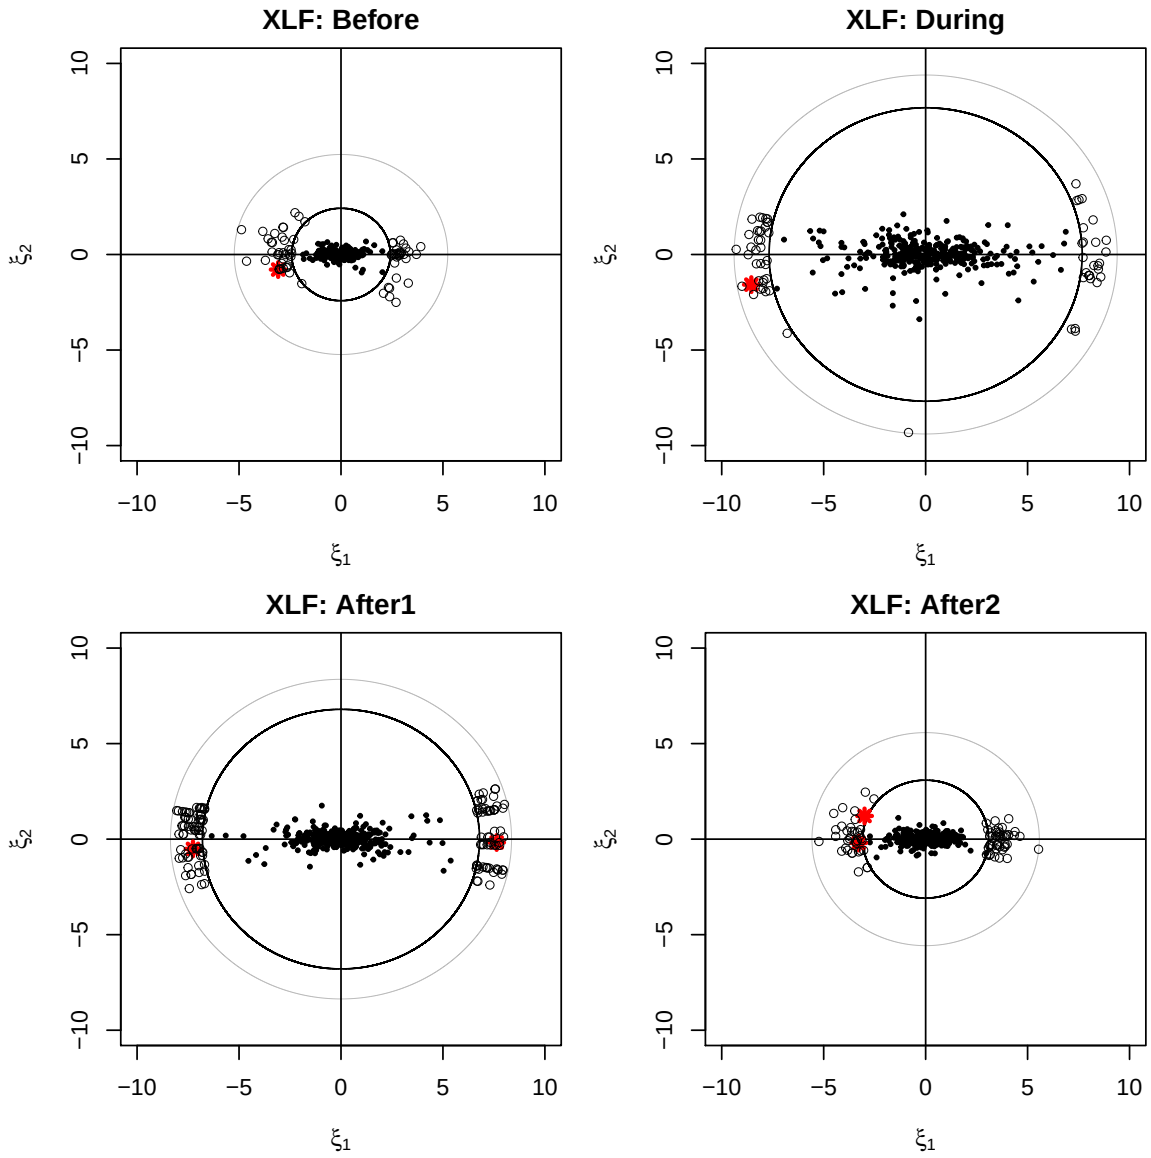

Figure D.3: **VaR Estimates (Bootstrap)**: The scatter plot of XLF scores (financials) with VaR estimates. On each plot, the big black thick circle corresponds to the norm at $\text{VaR}_{0.004}$ (or the 1-year return level), while the bigger grey circle corresponds to the norm at $\text{VaR}_{0.0001}$. The stars represent the real extreme points with norm greater than $\text{VaR}_{0.004}$. The little grey circles represent the simulated extreme points with norm greater than $\text{VaR}_{0.004}$ based on the Bootstrap method.

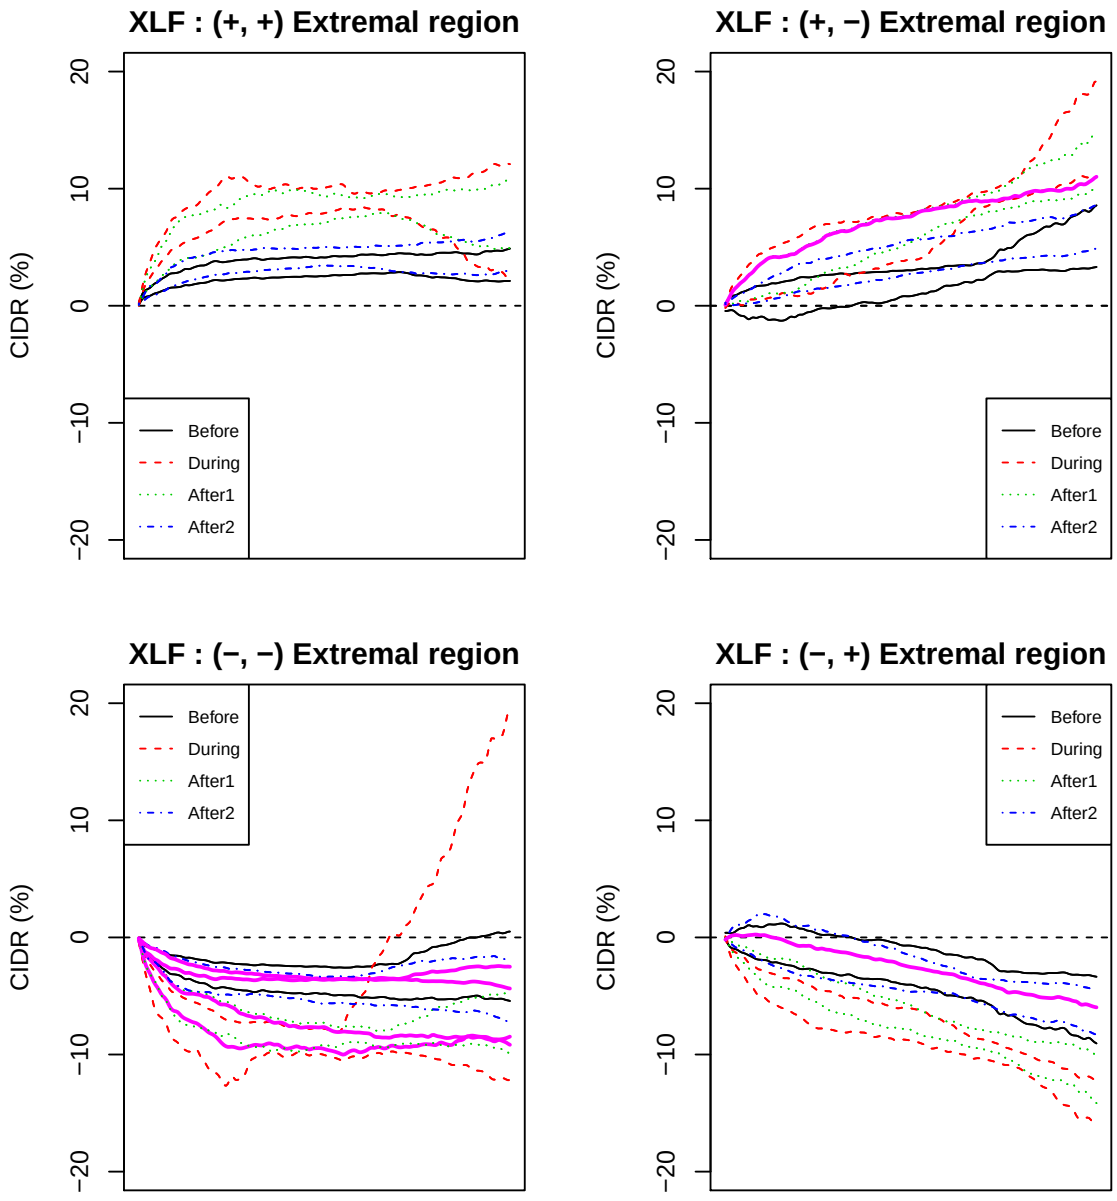

Figure D.4: **VaR Extremal regions (Bootstrap):** The extremal regions of CIDRs corresponding to extreme regions in different quadrants on the (ξ_1, ξ_2) plane in Figure D.3. On each panel, the two curves in the same type, respectively, represents the lower and upper bound of the extremal region of CIDRs regarding one of the four periods under comparison. The thick solid curves represent the real extremal CIDRs with norm greater than $\text{VaR}_{0.004}$.

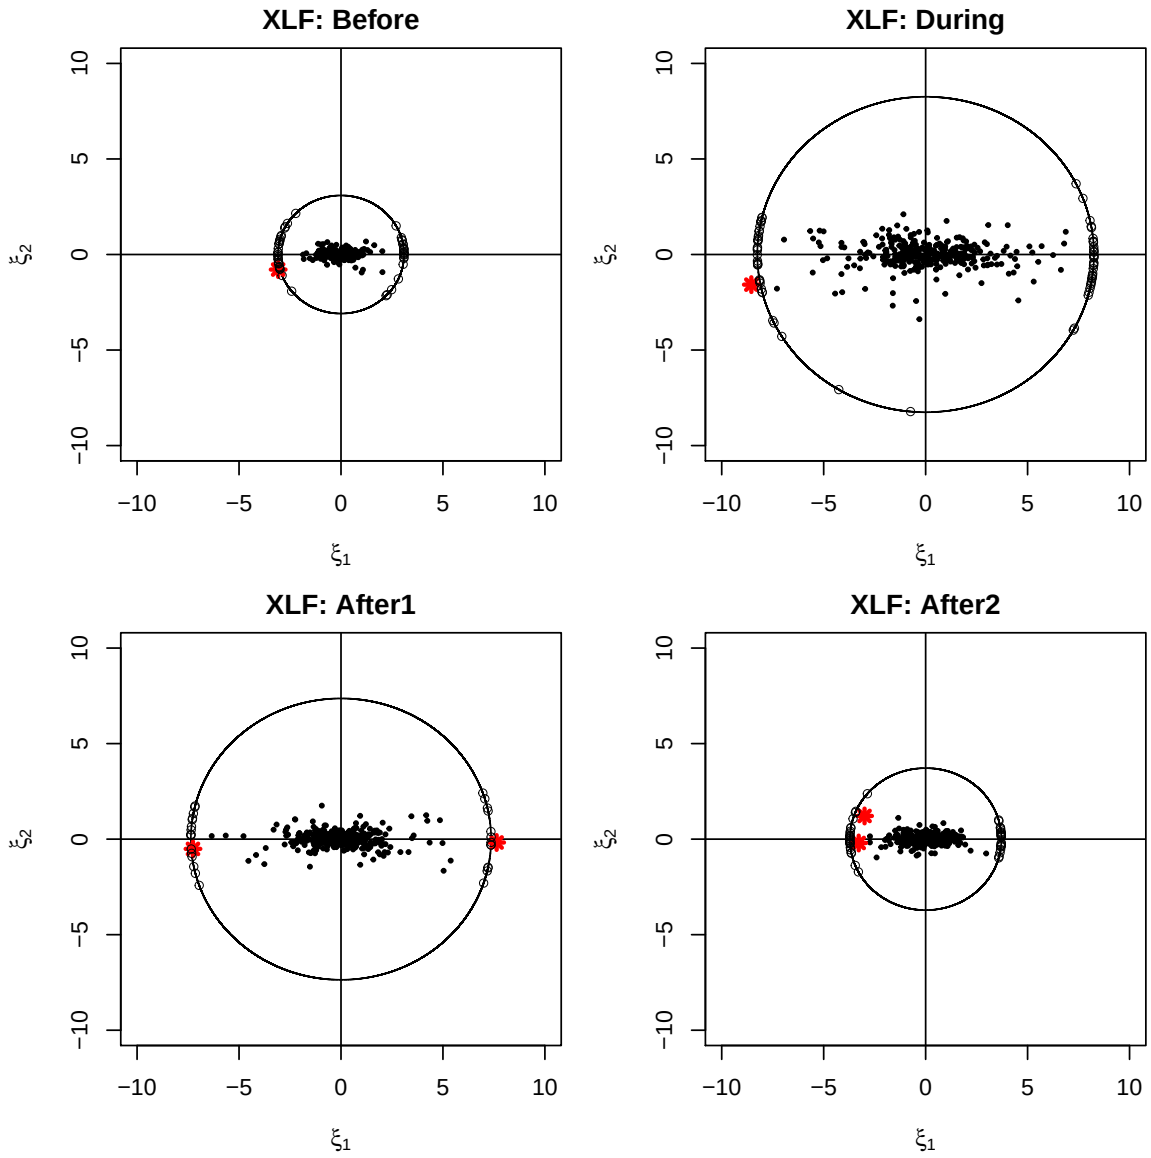

Figure D.5: **ES Estimates (Bootstrap)**: The scatter plot of XLF scores (financials) with ES estimates. On each plot, the big black thick circle corresponds to the norm at $\text{ES}_{0.004}$ (or the 1-year expected shortfall). The stars represent the real extreme points with norms greater than $\text{VaR}_{0.004}$. The little grey circles represent the simulated extreme points with norm at $\text{ES}_{0.004}$ based on the Bootstrap method.

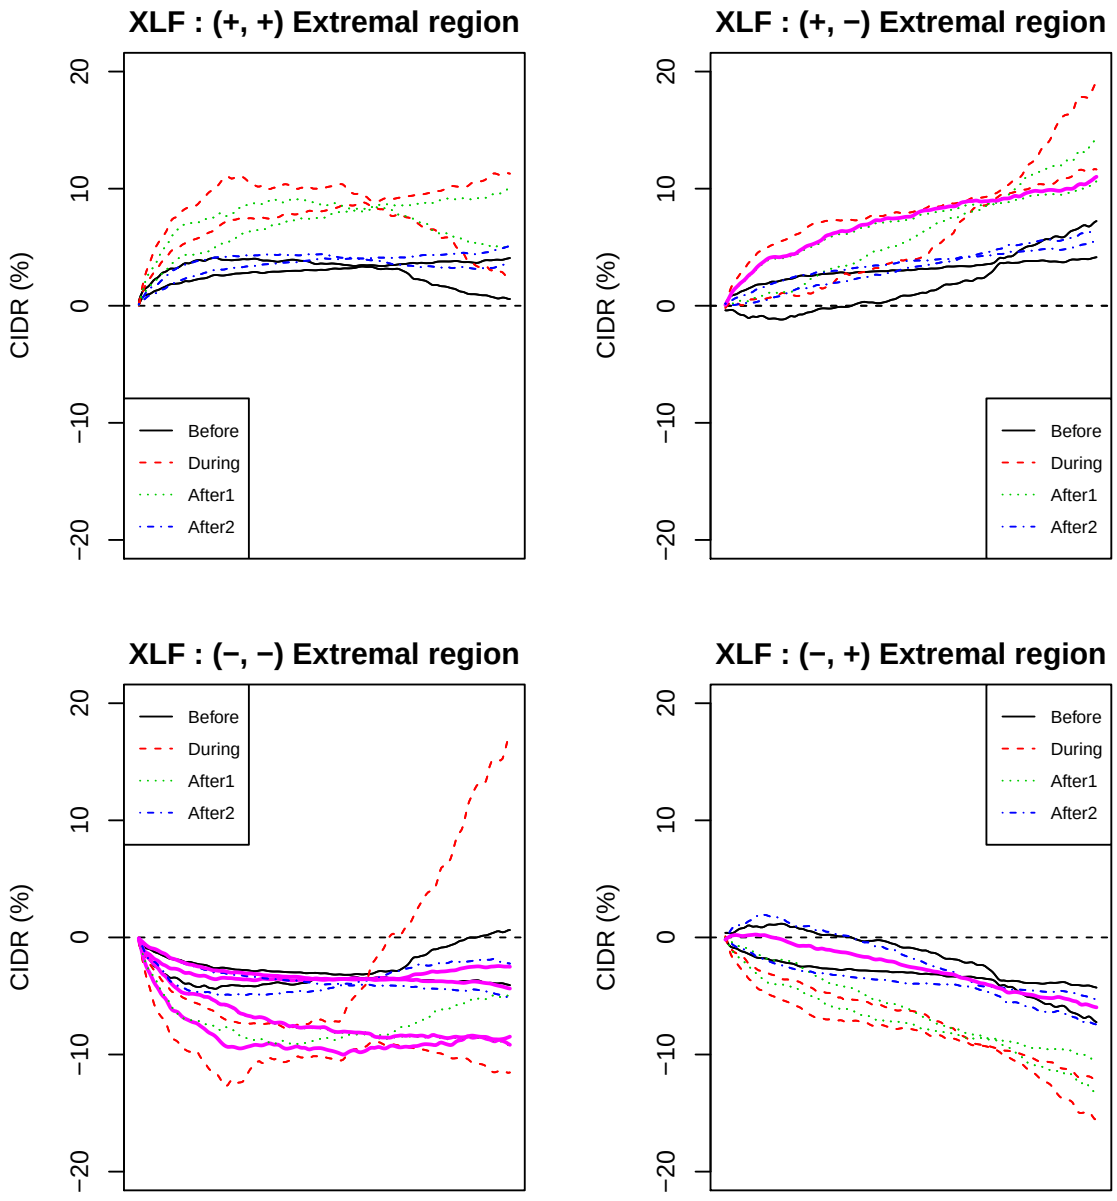

Figure D.6: **ES Extremal regions (Bootstrap):** The extremal regions of CIDRs corresponding to extreme regions in different quadrants on the (ξ_1, ξ_2) plane in Figure D.5. On each panel, the two curves in the same type, respectively, represents the lower and upper bound of the extremal region of CIDRs regarding one of the four periods under comparison. The thick solid curves represent the real extremal CIDRs with norm greater than $\text{VaR}_{0.004}$.

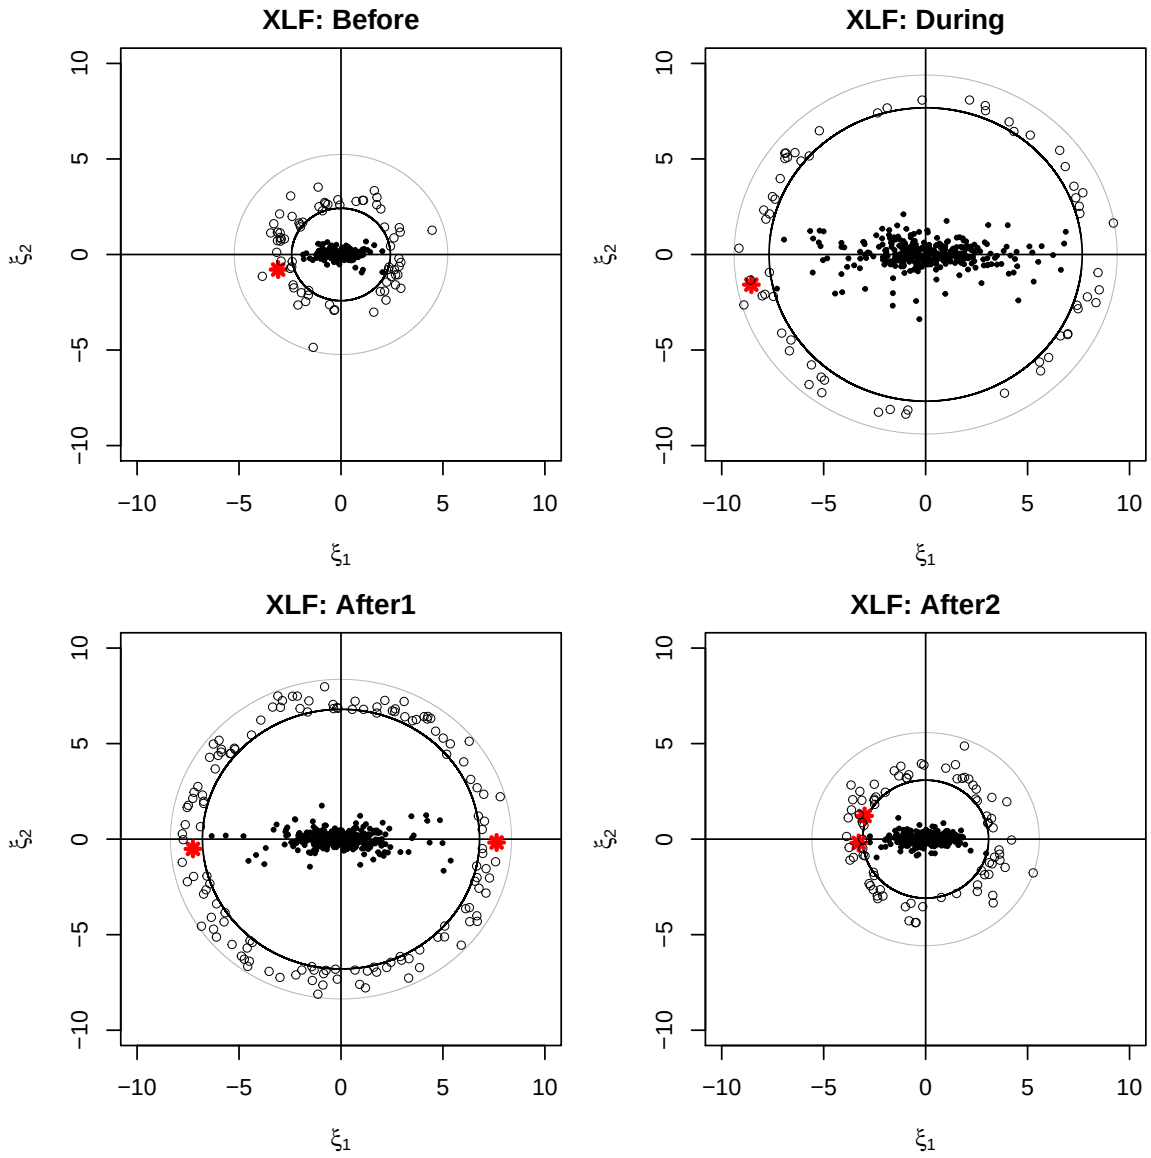

Figure D.7: **VaR Estimates (KDE)**: The scatter plot of XLF scores (financials) with VaR estimates. On each plot, the big black thick circle corresponds to the norm at $\text{VaR}_{0.004}$ (or the 1-year return level), while the bigger grey circle corresponds to the norm at $\text{VaR}_{0.0001}$. The stars represent the real extreme points with norm greater than $\text{VaR}_{0.004}$. The little grey circles represent the simulated extreme points with norm greater than $\text{VaR}_{0.004}$ based on KDE.

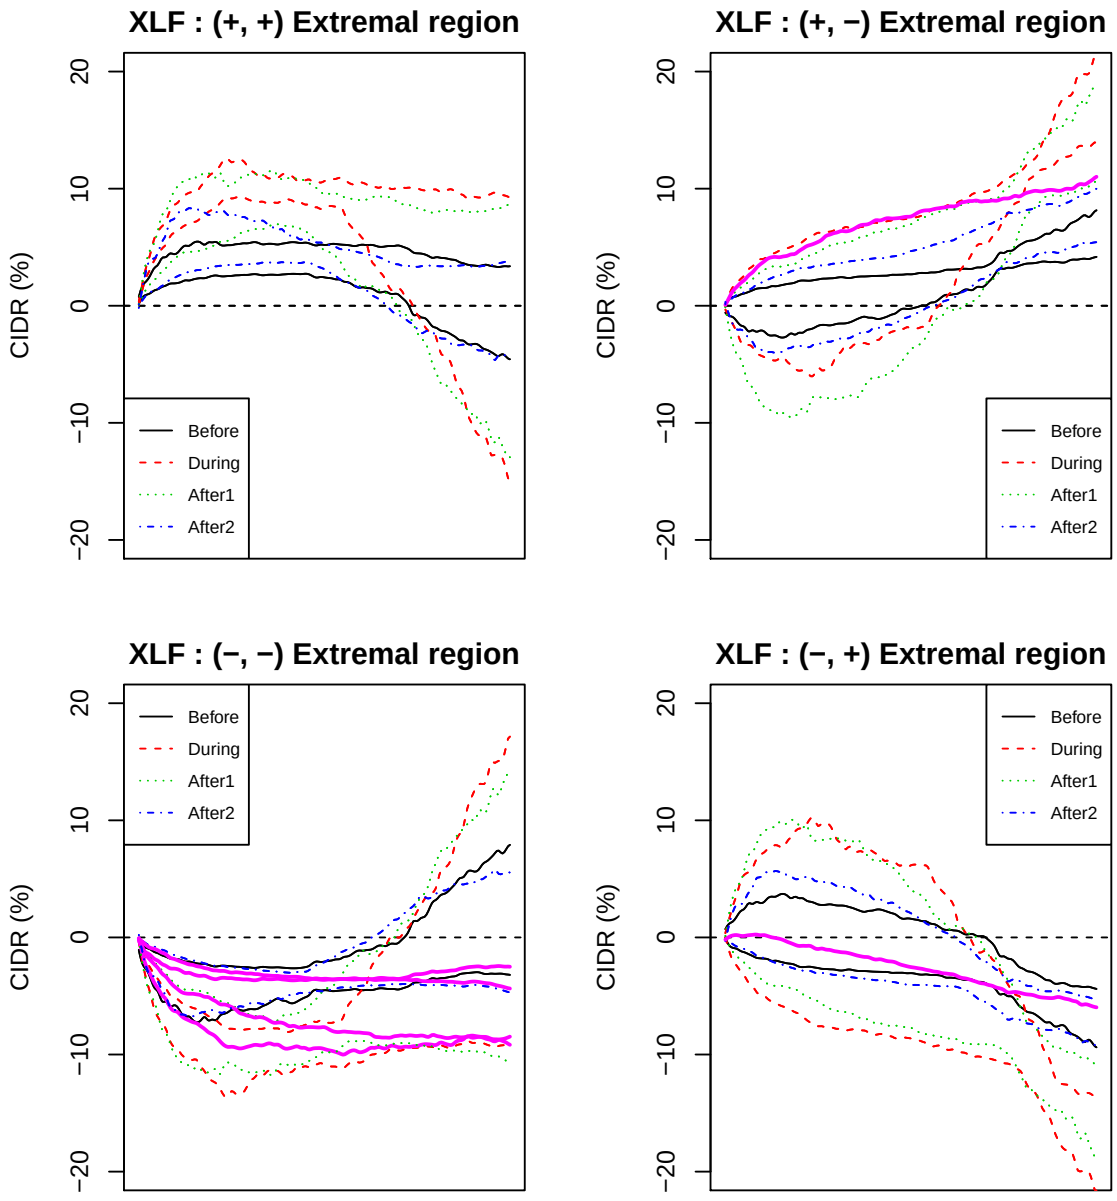

Figure D.8: **VaR Extremal regions (KDE):** The extremal regions of CIDRs corresponding to extreme regions in different quadrants on the (ξ_1, ξ_2) plane in Figure D.7. On each panel, the two curves in the same type, respectively, represents the lower and upper bound of the extremal region of CIDRs regarding one of the four periods under comparison. The thick solid curves represent the real extremal CIDRs with norm greater than $\text{VaR}_{0.004}$.

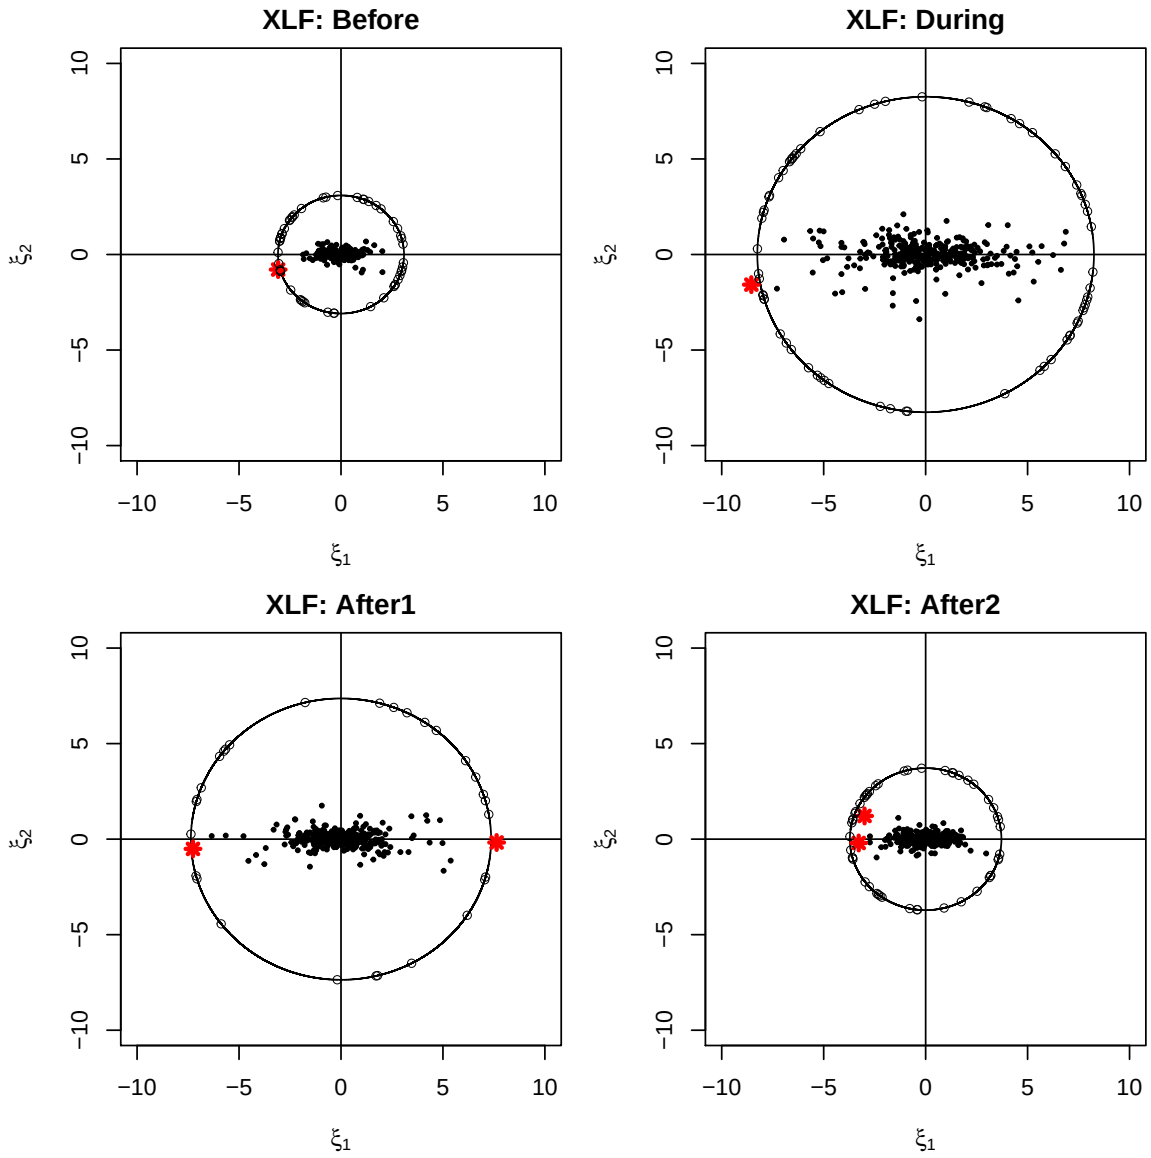

Figure D.9: **ES Estimates (KDE)**: The scatter plot of XLF scores (financials) with ES estimates. On each plot, the big black thick circle corresponds to the norm at $\text{ES}_{0.004}$ (or the 1-year expected shortfall). The stars represent the real extreme points with norms greater than $\text{VaR}_{0.004}$. The little grey circles represent the simulated extreme points with norm at $\text{ES}_{0.004}$ based on KDE.

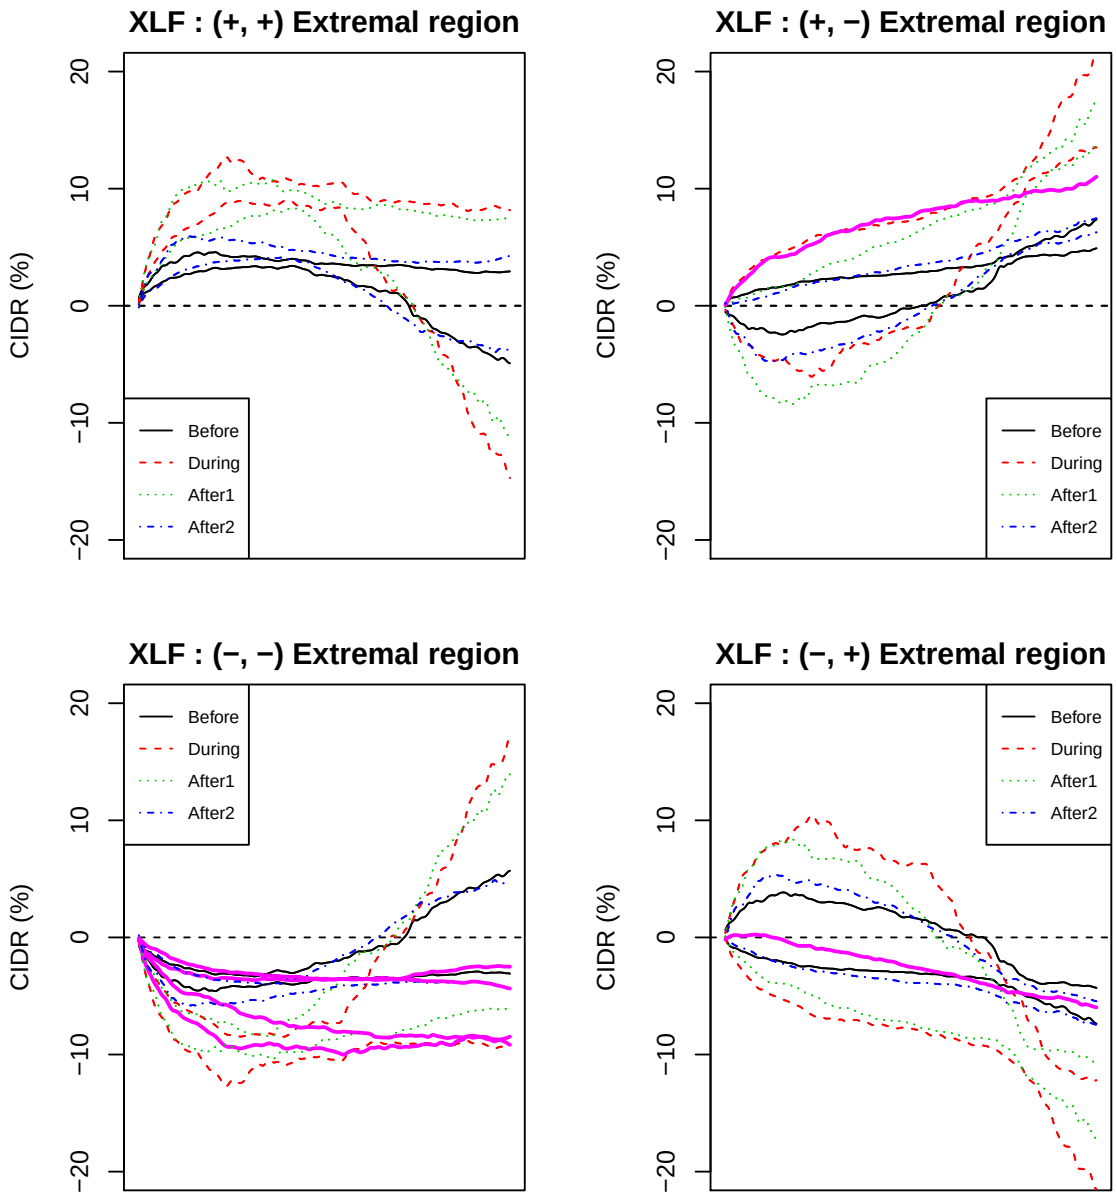

Figure D.10: **ES Extremal regions (KDE):** The extremal regions of CIDRs corresponding to extreme regions in different quadrants on the (ξ_1, ξ_2) plane in Figure D.9. On each panel, the two curves in the same type, respectively, represents the lower and upper bound of the extremal region of CIDRs regarding one of the four periods under comparison. The thick solid curves represent the real extremal CIDRs with norm greater than $\text{VaR}_{0.004}$.

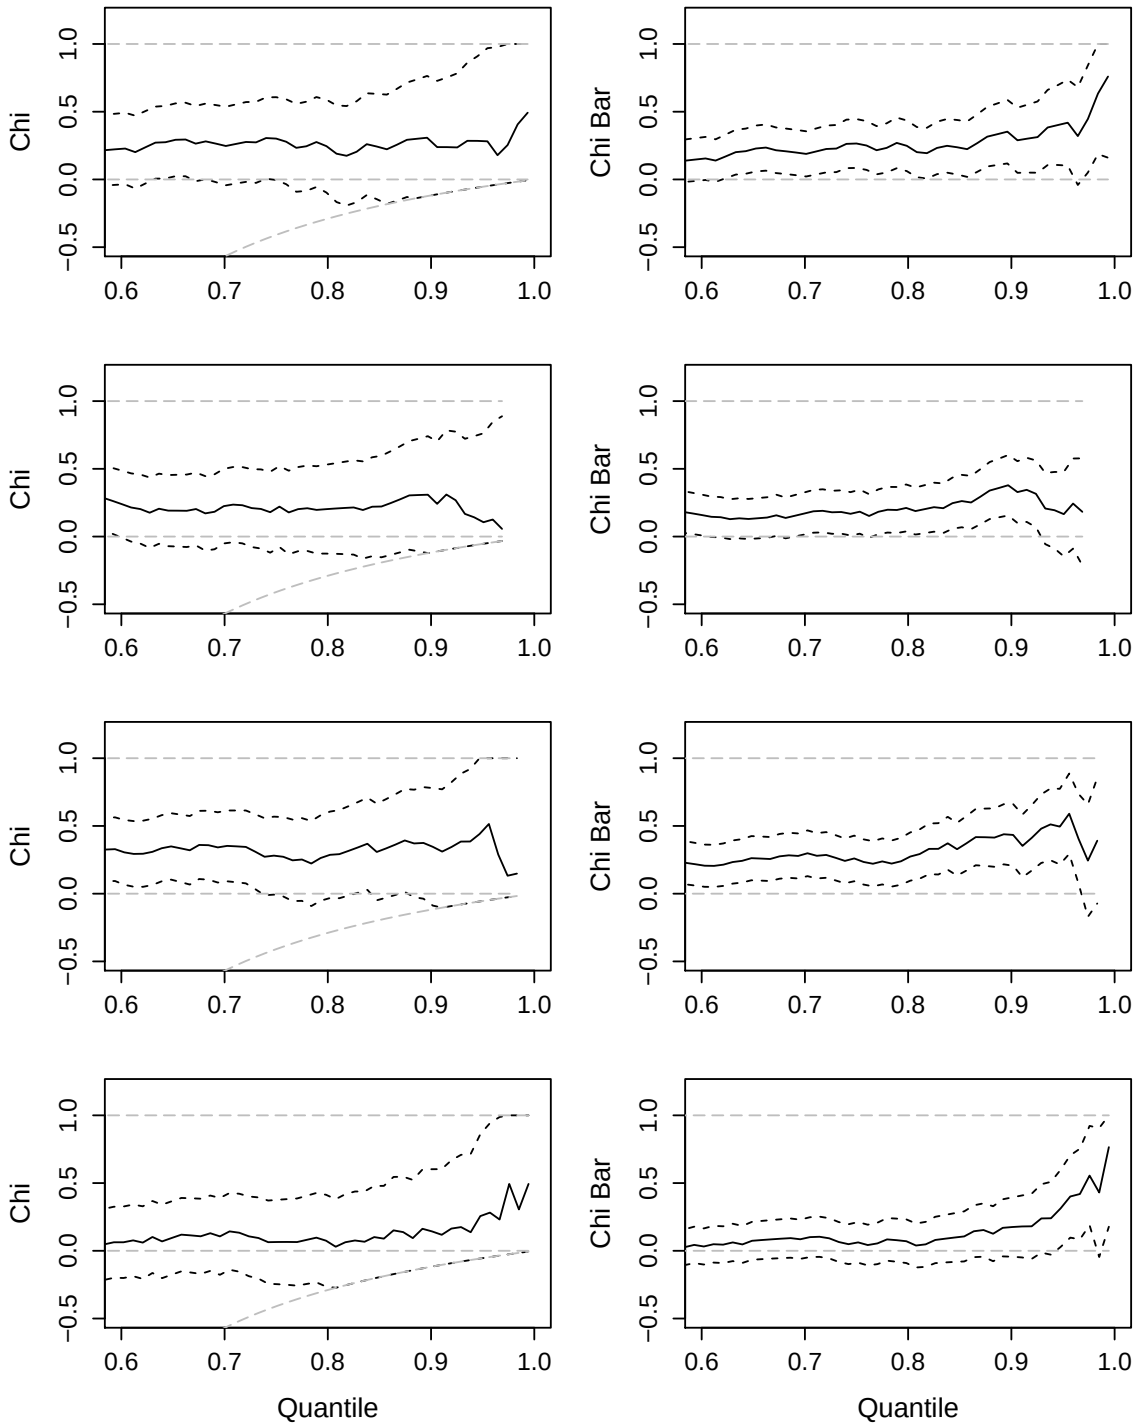

Figure D.11: Estimates of $\chi(q)$ and $\overline{\chi}(q)$ for the squared scores of the first two CIDR FPCs for XLF. Dotted lines correspond to approximate 95% confidence intervals. From top to bottom, the panels show the results for the periods "before", "during" "after1" and "after2".

ETF	Prd	u	N_u	Parameter $\hat{\gamma}$	$\text{VaR}_{.004}$	$\text{ES}_{.004}$	$\text{VaR}_{.0008}$	$\text{ES}_{.0008}$
	Bf	0.8	54	0.12(0.16)	2.46(0.16)	3.08(0.13)	3.43(0.94)	4.18(0.34)
XLF	Dr	2.7	71	-0.27(0.09)	7.83(0.20)	8.40(0.70)	8.78(0.43)	9.15(0.97)
(Financials)	A1	2.5	33	-0.24(0.20)	6.81(0.37)	7.54(1.57)	8.02(1.23)	8.52(2.56)
	A2	1.3	51	0.02(0.16)	3.11(0.13)	3.65(0.19)	3.98(0.69)	4.54(0.35)
	Bf	0.8	67	-0.08(0.11)	2.13(0.04)	2.39(0.09)	2.56(0.13)	2.79(0.17)
XLK	Dr	1.6	69	-0.08(0.11)	4.96(0.23)	5.64(0.43)	6.06(0.79)	6.66(0.74)
(Technology)	A1	1.1	73	0.16(0.12)	3.08(0.19)	3.90(0.19)	4.33(1.09)	5.39(0.65)
	A2	0.8	68	-0.10(0.15)	1.99(0.03)	2.23(0.11)	2.37(0.14)	2.57(0.21)
	Bf	0.6	78	0.09(0.11)	1.91(0.06)	2.33(0.05)	2.58(0.31)	3.06(0.12)
XLY	Dr	2.2	58	-0.02(0.13)	7.06(0.69)	8.29(1.11)	9.04(2.92)	10.22(2.03)
(Cons. Dis.)	A1	1.4	63	0.10(0.17)	4.20(0.41)	5.18(0.33)	5.74(2.49)	6.88(0.64)
	A2	0.9	69	-0.06(0.15)	2.56(0.08)	2.92(0.16)	3.14(0.35)	3.47(0.31)
	Bf	0.6	50	-0.15(0.15)	1.34(0.01)	1.47(0.08)	1.55(0.04)	1.65(0.15)
XLP	Dr	0.6	185	0.20(0.08)	5.38(0.82)	7.27(0.43)	8.25(4.41)	10.84(1.37)
(Cons. Stap.)	A1	0.8	72	-0.10(0.12)	1.92(0.02)	2.13(0.07)	2.27(0.09)	2.45(0.14)
	A2	0.6	66	0.12(0.14)	1.51(0.04)	1.84(0.06)	2.03(0.21)	2.43(0.21)
	Bf	1.3	50	-0.08(0.19)	2.46(0.05)	2.71(0.16)	2.87(0.23)	3.09(0.32)
XLE	Dr	2.5	51	-0.23(0.14)	5.79(0.16)	6.27(0.60)	6.58(0.49)	6.91(0.94)
(Energy)	A1	1.7	51	-0.23(0.15)	3.45(0.05)	3.71(0.23)	3.88(0.14)	4.06(0.39)
	A2	1.3	57	-0.12(0.13)	2.93(0.05)	3.24(0.16)	3.44(0.19)	3.70(0.30)
	Bf	0.5	66	0.11(0.14)	1.74(0.09)	2.24(0.18)	2.51(0.54)	3.16(0.77)
XLV	Dr	1.5	53	-0.06(0.17)	4.46(0.29)	5.15(0.58)	5.59(1.31)	6.22(1.08)
(Health)	A1	0.8	95	0.12(0.11)	3.13(0.20)	3.93(0.13)	4.38(1.01)	5.35(0.30)
	A2	0.7	71	0.11(0.11)	2.20(0.08)	2.73(0.08)	3.03(0.41)	3.66(0.19)
	Bf	0.8	48	-0.12(0.16)	1.69(0.02)	1.87(0.10)	1.98(0.08)	2.12(0.20)
XLI	Dr	1.5	88	0.14(0.13)	5.28(0.68)	6.67(0.41)	7.44(3.92)	9.17(0.88)
(Industrials)	A1	1.5	65	-0.10(0.11)	3.83(0.10)	4.28(0.22)	4.57(0.31)	4.95(0.39)
	A2	1.2	48	-0.10(0.17)	2.42(0.04)	2.69(0.14)	2.85(0.17)	3.08(0.28)
	Bf	0.8	85	0.09(0.13)	2.75(0.16)	3.35(0.09)	3.70(0.81)	4.39(0.18)
XLB	Dr	2.0	74	0.02(0.14)	5.83(0.51)	6.87(0.57)	7.49(2.40)	8.57(1.07)
(Materials)	A1	1.5	70	-0.05(0.14)	3.99(0.17)	4.54(0.30)	4.88(0.71)	5.38(0.54)
	A2	1.2	64	-0.23(0.10)	2.84(0.03)	3.06(0.13)	3.21(0.07)	3.36(0.21)
	Bf	0.8	57	-0.01(0.15)	2.17(0.07)	2.52(0.10)	2.74(0.32)	3.08(0.20)
XLU	Dr	0.9	163	0.13(0.09)	5.78(0.75)	7.35(0.34)	8.22(3.53)	10.13(0.67)
(Utilities)	A1	1.1	56	-0.06(0.14)	2.83(0.08)	3.24(0.17)	3.49(0.32)	3.86(0.32)
	A2	0.7	74	0.16(0.14)	2.17(0.12)	2.77(0.17)	3.10(0.74)	3.88(0.73)

Table D.1: Estimation results for the norm $\|R\|$ for nine Sector ETFs. The Value-at-Risk estimates $\text{VaR}_{.004}$ and $\text{VaR}_{.0008}$, respectively, represent the 1-year and 5-years return levels. Standard errors in parentheses. Recall that u is the threshold and N_u is the sample size of values over the threshold.

ETF	Prd	u	N_u	Parameter $\hat{\gamma}$	$\text{VaR}_{.004}$	$\text{ES}_{.004}$	$\text{VaR}_{.0008}$	$\text{ES}_{.0008}$
	Bf	0.7	60	0.10(0.14)	2.30(0.12)	2.86(0.10)	3.17(0.64)	3.82(0.21)
XLF	Dr	2.5	69	-0.30(0.10)	7.57(0.18)	8.11(0.69)	8.46(0.38)	8.79(0.95)
(Financials)	A1	2.3	34	-0.23(0.19)	6.72(0.40)	7.49(1.58)	8.00(1.35)	8.54(2.60)
	A2	1.2	57	0.02(0.15)	2.98(0.12)	3.48(0.16)	3.78(0.58)	4.30(0.30)
	Bf	0.9	49	-0.07(0.15)	2.07(0.04)	2.34(0.10)	2.50(0.16)	2.74(0.20)
XLK	Dr	1.6	60	-0.06(0.13)	4.91(0.29)	5.66(0.52)	6.13(1.14)	6.82(0.94)
(Technology)	A1	1.0	76	0.14(0.12)	3.05(0.18)	3.84(0.16)	4.27(0.99)	5.25(0.43)
	A2	0.9	45	-0.03(0.23)	1.97(0.06)	2.26(0.13)	2.44(0.33)	2.72(0.27)
	Bf	0.6	68	0.12(0.13)	1.87(0.08)	2.32(0.07)	2.58(0.40)	3.12(0.22)
XLY	Dr	2.2	51	-0.04(0.14)	6.97(0.64)	8.14(1.17)	8.86(2.56)	9.95(2.13)
(Cons. Dis.)	A1	1.0	108	0.06(0.11)	4.04(0.28)	4.88(0.21)	5.37(1.23)	6.30(0.37)
	A2	1.0	48	-0.14(0.17)	2.45(0.05)	2.74(0.20)	2.92(0.21)	3.14(0.39)
	Bf	0.6	44	-0.19(0.14)	1.31(0.01)	1.43(0.07)	1.50(0.03)	1.59(0.15)
XLP	Dr	0.5	188	0.19(0.08)	5.22(0.77)	7.06(0.40)	8.03(4.06)	10.54(1.24)
(Cons. Stap.)	A1	0.9	48	-0.07(0.15)	1.89(0.03)	2.13(0.08)	2.28(0.12)	2.49(0.17)
	A2	0.5	90	0.13(0.12)	1.47(0.04)	1.81(0.07)	2.01(0.21)	2.43(0.29)
	Bf	1.2	52	-0.21(0.15)	2.39(0.02)	2.57(0.15)	2.68(0.08)	2.80(0.26)
XLE	Dr	2.5	45	-0.14(0.17)	5.62(0.24)	6.24(0.71)	6.64(0.96)	7.14(1.26)
(Energy)	A1	1.7	46	-0.23(0.16)	3.40(0.05)	3.66(0.24)	3.83(0.14)	4.01(0.42)
	A2	1.2	57	-0.17(0.12)	2.85(0.04)	3.13(0.15)	3.31(0.13)	3.52(0.27)
	Bf	0.5	58	0.15(0.16)	1.69(0.09)	2.17(0.16)	2.43(0.53)	3.05(0.66)
XLV	Dr	1.7	35	-0.18(0.22)	4.31(0.18)	4.81(0.80)	5.14(0.79)	5.51(1.45)
(Health)	A1	0.7	104	0.10(0.10)	3.07(0.18)	3.83(0.12)	4.26(0.84)	5.16(0.24)
	A2	0.7	59	0.08(0.11)	2.16(0.08)	2.66(0.08)	2.94(0.35)	3.51(0.16)
	Bf	0.8	39	-0.25(0.15)	1.64(0.01)	1.76(0.10)	1.84(0.03)	1.93(0.19)
XLI	Dr	1.3	88	0.11(0.12)	5.14(0.60)	6.43(0.42)	7.16(3.15)	8.70(0.82)
(Industrials)	A1	1.6	53	-0.07(0.13)	3.80(0.12)	4.30(0.26)	4.62(0.45)	5.07(0.48)
	A2	1.2	41	-0.09(0.21)	2.37(0.05)	2.65(0.18)	2.82(0.24)	3.06(0.38)
	Bf	0.8	76	0.13(0.15)	2.74(0.21)	3.44(0.13)	3.83(1.20)	4.68(0.38)
XLB	Dr	1.5	94	-0.04(0.11)	5.65(0.39)	6.52(0.52)	7.06(1.51)	7.86(0.91)
(Materials)	A1	1.6	48	-0.22(0.11)	3.94(0.07)	4.31(0.27)	4.55(0.17)	4.81(0.43)
	A2	1.3	50	-0.20(0.12)	2.79(0.03)	3.03(0.15)	3.19(0.09)	3.36(0.26)
	Bf	0.8	50	0.06(0.19)	2.16(0.11)	2.60(0.09)	2.86(0.62)	3.34(0.16)
XLU	Dr	0.9	136	0.06(0.08)	5.35(0.48)	6.56(0.34)	7.26(1.97)	8.60(0.61)
(Utilities)	A1	1.1	50	-0.01(0.16)	2.80(0.10)	3.27(0.18)	3.55(0.48)	4.00(0.35)
	A2	0.6	93	0.11(0.12)	2.04(0.08)	2.53(0.06)	2.81(0.41)	3.39(0.17)

Table D.2: Estimation results for the magnitude of the first score $|\xi_1|$ for nine Sector ETFs; quantities reported are the same as in Table D.1.

ETF	Prd	u	N_u	Parameter $\hat{\gamma}$	$\text{VaR}_{.004}$	$\text{ES}_{.004}$	$\text{VaR}_{.0008}$	$\text{ES}_{.0008}$
	Bf	1.1	59	0.06(0.18)	4.58(0.65)	5.72(0.47)	6.28(3.61)	7.53(0.97)
XLF	Dr	5.3	53	-0.23(0.13)	15.53(1.37)	16.00(2.77)	17.95(3.82)	17.96(4.27)
(Financials)	A1	4.9	26	-0.21(0.23)	13.52(2.03)	14.46(4.77)	16.47(7.46)	16.89(8.80)
	A2	2.4	53	-0.02(0.14)	7.77(0.90)	9.15(1.08)	10.06(3.94)	11.40(2.16)
	Bf	1.4	40	-0.03(0.14)	2.83(0.07)	3.16(0.08)	3.44(0.25)	3.76(0.17)
XLK	Dr	2.9	58	0.10(0.14)	9.20(1.77)	11.75(1.59)	12.72(9.59)	15.66(3.46)
(Technology)	A1	2.1	46	-0.10(0.12)	5.05(0.19)	5.51(0.30)	6.11(0.61)	6.48(0.57)
	A2	1.8	48	0.02(0.18)	4.89(0.41)	5.84(0.40)	6.37(2.17)	7.34(0.85)
	Bf	1.1	57	-0.04(0.15)	3.26(0.15)	3.71(0.17)	4.06(0.61)	4.48(0.35)
XLY	Dr	3.4	50	-0.04(0.14)	9.50(1.08)	10.90(1.33)	11.98(4.38)	13.29(2.61)
(Cons. Dis.)	A1	2.5	43	-0.15(0.12)	5.95(0.23)	6.31(0.42)	7.07(0.69)	7.28(0.76)
	A2	2.0	42	-0.05(0.16)	5.35(0.35)	6.11(0.50)	6.75(1.49)	7.43(1.02)
	Bf	0.8	51	0.04(0.14)	2.27(0.09)	2.76(0.08)	2.99(0.41)	3.50(0.16)
XLP	Dr	1.8	56	0.11(0.15)	6.65(1.11)	8.64(1.01)	9.43(6.23)	11.75(2.19)
(Cons. Stap.)	A1	1.2	60	-0.10(0.11)	3.24(0.08)	3.54(0.13)	3.91(0.24)	4.16(0.24)
	A2	1.2	42	0.00(0.16)	3.25(0.16)	3.87(0.19)	4.24(0.78)	4.86(0.41)
	Bf	1.7	67	-0.17(0.17)	3.82(0.11)	3.89(0.17)	4.34(0.41)	4.33(0.35)
XLE	Dr	3.5	64	0.04(0.14)	15.68(5.60)	19.38(5.57)	21.36(27.57)	25.29(11.57)
(Energy)	A1	2.7	51	-0.12(0.10)	6.16(0.21)	6.57(0.34)	7.30(0.58)	7.59(0.60)
	A2	2.3	57	-0.03(0.12)	6.90(0.57)	8.00(0.68)	8.77(2.24)	9.82(1.32)
	Bf	0.7	64	-0.12(0.15)	2.53(0.08)	2.78(0.17)	3.06(0.31)	3.25(0.32)
XLV	Dr	2.2	47	0.06(0.16)	7.84(1.47)	9.91(1.56)	10.88(7.99)	13.16(3.41)
(Health)	A1	1.6	38	-0.16(0.14)	3.98(0.12)	4.25(0.26)	4.77(0.37)	4.93(0.50)
	A2	1.4	57	-0.02(0.14)	4.44(0.29)	5.17(0.33)	5.68(1.26)	6.38(0.67)
	Bf	1.2	46	-0.10(0.18)	3.04(0.11)	3.33(0.17)	3.68(0.45)	3.92(0.36)
XLI	Dr	2.9	55	-0.17(0.12)	8.21(0.45)	8.67(0.81)	9.67(1.35)	9.92(1.36)
(Industrials)	A1	2.3	58	-0.10(0.13)	6.42(0.35)	7.03(0.51)	7.77(1.25)	8.25(0.94)
	A2	2.2	45	-0.12(0.14)	5.87(0.31)	6.41(0.52)	7.13(1.09)	7.55(0.98)
	Bf	1.3	59	-0.19(0.13)	3.72(0.09)	3.87(0.21)	4.31(0.28)	4.36(0.37)
XLB	Dr	4.1	40	-0.22(0.14)	10.59(0.70)	10.93(1.48)	12.37(1.98)	12.39(2.46)
(Materials)	A1	2.6	58	-0.32(0.12)	5.51(0.07)	5.19(0.20)	6.03(0.16)	5.58(0.32)
	A2	2.6	40	-0.09(0.19)	6.66(0.53)	7.40(0.77)	8.22(2.40)	8.83(1.60)
	Bf	1.3	43	-0.13(0.21)	3.80(0.21)	4.15(0.39)	4.62(0.95)	4.88(0.82)
XLU	Dr	2.4	51	0.11(0.17)	9.60(2.84)	12.66(2.66)	13.91(17.54)	17.51(6.12)
(Utilities)	A1	1.7	51	-0.25(0.15)	3.59(0.05)	3.52(0.14)	4.03(0.16)	3.86(0.27)
	A2	1.1	70	0.19(0.14)	3.90(0.44)	5.40(0.58)	5.81(2.87)	7.75(1.78)

Table D.3: Estimation results for the magnitude of the point-to-point returns for nine Sector ETFs; quantities reported are the same as in Table D.1.

ETF	Prd	u	N_u	(+, +)	(+, -)	(-, -)	(-, +)	P value
	Bf	0.8	52	15	8	15	14	0.27
XLF	Dr	2.7	68	13	21	18	16	0.62
(Financials)	A1	2.5	29	6	7	6	10	0.47
	A2	1.3	50	14	10	12	14	0.78
	Bf	0.8	61	14	12	15	20	1.00
XLK	Dr	1.6	64	18	15	15	16	1.00
(Technology)	A1	1.1	69	16	14	22	17	0.47
	A2	0.8	66	16	13	20	17	0.62
	Bf	0.6	74	19	16	18	21	1.00
XLY	Dr	2.2	53	12	16	12	13	0.59
(Cons. Dis.)	A1	1.4	61	18	12	16	15	0.44
	A2	0.9	63	18	15	10	20	0.44
	Bf	0.6	47	13	8	13	13	0.56
XLP	Dr	0.6	171	45	48	37	41	0.65
(Cons. Stap.)	A1	0.8	67	14	17	16	20	0.47
	A2	0.6	60	15	12	14	19	1.00
	Bf	1.3	46	13	12	8	13	0.56
XLE	Dr	2.5	48	12	11	8	17	0.38
(Energy)	A1	1.7	50	12	15	13	10	1.00
	A2	1.3	54	16	8	15	15	0.27
	Bf	0.5	63	15	17	16	15	1.00
XLV	Dr	1.5	51	16	8	13	14	0.39
(Health)	A1	0.8	92	25	19	29	19	0.14
	A2	0.7	65	17	10	19	19	0.32
	Bf	0.8	44	10	11	12	11	1.00
XLI	Dr	1.5	77	21	18	17	21	1.00
(Industrials)	A1	1.5	63	16	15	17	15	0.80
	A2	1.2	46	10	11	14	11	1.00
	Bf	0.8	81	21	19	20	21	1.00
XLB	Dr	2.0	67	19	12	11	25	0.61
(Materials)	A1	1.5	60	13	15	17	15	1.00
	A2	1.2	61	20	11	13	17	0.61
	Bf	0.8	55	15	12	15	13	0.59
XLU	Dr	0.9	147	39	38	34	36	1.00
(Utilities)	A1	1.1	53	11	13	14	15	0.78
	A2	0.7	69	20	18	13	18	0.81

Table D.4: Fisher's exact test on independence of the four regions with frequencies of points falling in each region given their norms above the threshold.

ETF	Prd	u	N_u	P(+, +)	P(+, -)	P(-, -)	P(-, +)	P value
	Bf	0.8	52	0.29	0.15	0.29	0.27	0.27
XLF	Dr	2.7	68	0.19	0.31	0.26	0.24	0.62
(Financials)	A1	2.5	29	0.21	0.24	0.21	0.34	0.47
	A2	1.3	50	0.28	0.20	0.24	0.28	0.78
	Bf	0.8	61	0.23	0.20	0.25	0.33	1.00
XLK	Dr	1.6	64	0.28	0.23	0.23	0.25	1.00
(Technology)	A1	1.1	69	0.23	0.20	0.32	0.25	0.47
	A2	0.8	66	0.24	0.20	0.30	0.26	0.62
	Bf	0.6	74	0.26	0.22	0.24	0.28	1.00
XLY	Dr	2.2	53	0.23	0.30	0.23	0.25	0.59
(Cons. Dis.)	A1	1.4	61	0.30	0.20	0.26	0.25	0.44
	A2	0.9	63	0.29	0.24	0.16	0.32	0.44
	Bf	0.6	47	0.28	0.17	0.28	0.28	0.56
XLP	Dr	0.6	171	0.26	0.28	0.22	0.24	0.65
(Cons. Stap.)	A1	0.8	67	0.21	0.25	0.24	0.30	0.47
	A2	0.6	60	0.25	0.20	0.23	0.32	1.00
	Bf	1.3	46	0.28	0.26	0.17	0.28	0.56
XLE	Dr	2.5	48	0.25	0.23	0.17	0.35	0.38
(Energy)	A1	1.7	50	0.24	0.30	0.26	0.20	1.00
	A2	1.3	54	0.30	0.15	0.28	0.28	0.27
	Bf	0.5	63	0.24	0.27	0.25	0.24	1.00
XLV	Dr	1.5	51	0.31	0.16	0.25	0.27	0.39
(Health)	A1	0.8	92	0.27	0.21	0.32	0.21	0.14
	A2	0.7	65	0.26	0.15	0.29	0.29	0.32
	Bf	0.8	44	0.23	0.25	0.27	0.25	1.00
XLI	Dr	1.5	77	0.27	0.23	0.22	0.27	1.00
(Industrials)	A1	1.5	63	0.25	0.24	0.27	0.24	0.80
	A2	1.2	46	0.22	0.24	0.30	0.24	1.00
	Bf	0.8	81	0.26	0.23	0.25	0.26	1.00
XLB	Dr	2.0	67	0.28	0.18	0.16	0.37	0.61
(Materials)	A1	1.5	60	0.22	0.25	0.28	0.25	1.00
	A2	1.2	61	0.33	0.18	0.21	0.28	0.61
	Bf	0.8	55	0.27	0.22	0.27	0.24	0.59
XLU	Dr	0.9	147	0.27	0.26	0.23	0.24	1.00
(Utilities)	A1	1.1	53	0.21	0.25	0.26	0.28	0.78
	A2	0.7	69	0.29	0.26	0.19	0.26	0.81

Table D.5: Fisher's exact test on independence of the four regions with probabilities of points falling in each region given their norms above the threshold.

ETF	Prd	q	n_u	$\widehat{\overline{\chi}}$	Pvalue	$\widehat{\chi}$
	Bf	0.80	63	0.59(0.43)	0.21*	0.40(0.04)
XLF	Dr	0.70	106	0.25(0.32)	0.00	0
(Financials)	A1	0.80	71	0.53(0.42)	0.15*	0.43(0.05)
	A2	0.80	72	0.73(0.41)	0.39*	0.35(0.04)
	Bf	0.80	63	0.67(0.43)	0.32*	0.38(0.04)
XLK	Dr	0.70	105	0.01(0.33)	0.00	0
(Technology)	A1	0.80	71	-0.20(0.35)	0.00	0
	A2	0.75	89	0.21(0.33)	0.00	0
	Bf	0.70	94	0.62(0.34)	0.15*	0.42(0.04)
XLY	Dr	0.75	87	0.28(0.39)	0.02	0
(Cons. Dis.)	A1	0.65	125	0.31(0.28)	0.00	0
	A2	0.75	90	0.37(0.37)	0.04	0
	Bf	0.80	63	0.29(0.41)	0.04	0
XLP	Dr	0.60	140	0.62(0.33)	0.13*	0.42(0.03)
(Cons. Stap.)	A1	0.85	54	-0.03(0.43)	0.01	0
	A2	0.75	90	0.42(0.38)	0.05	0
	Bf	0.80	63	-0.59(0.32)	0.00	0
XLE	Dr	0.80	70	0.15(0.40)	0.01	0
(Energy)	A1	0.80	71	0.21(0.37)	0.01	0
	A2	0.70	108	0.61(0.35)	0.14*	0.39(0.03)
	Bf	0.80	63	0.97(0.50)	0.93*	0.38(0.04)
XLV	Dr	0.80	70	-0.10(0.33)	0.00	0
(Health)	A1	0.70	106	0.25(0.30)	0.00	0
	A2	0.85	54	0.19(0.42)	0.02	0
	Bf	0.85	47	0.54(0.47)	0.22*	0.37(0.05)
XLI	Dr	0.80	70	-0.58(0.28)	0.00	0
(Industrials)	A1	0.80	70	-0.16(0.35)	0.00	0
	A2	0.80	72	0.23(0.39)	0.02	0
	Bf	0.85	47	0.44(0.50)	0.16*	0.35(0.05)
XLB	Dr	0.80	70	-0.37(0.34)	0.00	0
(Materials)	A1	0.80	71	-0.36(0.30)	0.00	0
	A2	0.80	72	0.13(0.36)	0.01	0
	Bf	0.80	62	0.07(0.37)	0.00	0
XLU	Dr	0.80	70	0.09(0.38)	0.01	0
(Utilities)	A1	0.80	71	0.31(0.40)	0.04	0
	A2	0.80	72	0.87(0.45)	0.70*	0.36(0.04)

Table D.6: Estimation results of extreme dependence between ξ_1^2 and ξ_2^2 for the nine ETF sectors. Standard errors in parentheses. The value of q is the quantile for the threshold. Pvalue is for the likelihood ratio test of $\overline{\chi} = 1$, and "*" indicates cases where $\overline{\chi}$ is not significantly different from 1.