Master Degree Programme
in Data Analytics for Business and Society

Final Thesis

Modeling Extremal Events in Cryptocurrency and Stablecoin Markets Using Univariate and Bivariate Extreme Value Theory

Supervisor
Ch.ma Prof.ssa Ilaria Prosdocimi

Graduand
Lisa Adami
Matriculation Number 905259

Academic Year
2025 / 2026

Contents

Declaration on the Use of Artificial Intelligence (AI) Tools

The undersigned **Lisa Adami**, student ID number **905259**, enrolled at Ca' Foscari in the Master's Degree Programme in **DABS**, author of the degree thesis **"Modeling Extremal Events in Cryptocurrency and Stablecoin Markets Using Univariate and Bivariate Extreme Value Theory"**, thesis supervisor **Ilaria Prosdocimi**, aware of the penalties provided for in Article 76 of the Consolidated Act set out in Presidential Decree no. 445 of 28 December 2000, and of the loss of benefits provided for in Article 75 of the same Consolidated Act in the event of false or misleading statements, under her own responsibility pursuant to Articles 46 and 47 of Presidential Decree no. 445 of 28 December 2000,

DECLARES

- ☐ that she has NOT used generative artificial intelligence technologies to support various aspects of the research and writing of the final paper/degree thesis;

- ☑ that she HAS used generative artificial intelligence technologies to support various aspects of the research and writing of his/her final paper/degree thesis, in particular (for each point, write in full the requested information):

  - indicate the name of the tool (e.g., “ChatGPT” if an application or “GPT” if a model), with reference to the specific version (for example, “ChatGPT-4, version released in March 2024”);

    Gemini 3.5 Flash (unive), Claude Sonnet 4.6

  - if more than one tool has been used, specify each of them by indicating its name and version;

    Gemini 3.5 Flash (unive), Claude Sonnet 4.6

  - provide information on how the tool was used:
    - ☑ revision and optimisation of texts to improve clarity and coherence;
    - ☐ generation of charts and diagrams for data presentation;
    - ☐ support for the automatic drafting of sentences;
    - ☐ production of illustrative images to complement the methodological section;
    - ☑ other (please specify); **support for the resolution of errors in LaTeX code**

  - indicate the sections or chapters concerned;

    Linguistic revision concerned the entire thesis, with particular attention to the more discursive parts (Chapters 1, 7 and 8). Support for error resolution mainly concerned the correct formatting of citations and references, as well as the consistency of table and figure formatting.

  - indicate how the accuracy, quality and reliability of any information or text that may have been generated were verified.

    Linguistic revision was carried out through the use of online dictionaries and re-reading of the produced text. LaTeX code was verified through compilation and by checking formatting consistency and the absence of errors.

DECLARES FURTHER

that all material generated by AI has been carefully checked, reviewed and reworked by the author, who assumes full responsibility for the quality, accuracy and scientific integrity of the work.

Adami Lisa

Abstract

Cryptocurrency markets are characterized by large price fluctuations and recurrent episodes of market stress, making the assessment of tail risk particularly important. Standard Gaussian-based models often underestimate the probability of extreme events because cryptocurrency returns exhibit heavy tails and volatility clustering.

This thesis investigates the tail behavior of four cryptocurrencies: *Bitcoin* (BTC), *Ethereum* (ETH), *Tether* (USDT) and *USD Coin* (USDC), using Extreme Value Theory (EVT) to compare the risk characteristics of highly volatile assets and stablecoins within a unified framework.

Return series were filtered using GARCH-type models, with ARMA-GARCH specifications adopted for the stablecoins in order to satisfy the assumptions required by EVT. The resulting standardized residuals were analyzed using two univariate EVT approaches: the Block Maxima Method (BMM), based on the Generalized Extreme Value (GEV) distribution, and the Peaks Over Threshold (POT) method, based on the Generalized Pareto Distribution (GPD). Return levels were estimated over a five-year horizon and used as the primary measure of tail risk.

The results reveal substantial differences between volatile cryptocurrencies and stablecoins. Bitcoin and Ethereum exhibit considerably larger extreme gains and losses than USDT and USDC. While the stablecoins display Fréchet-type tail behavior, the filtered residuals of Bitcoin and Ethereum do not provide sufficient evidence to reject a Gumbel-type specification. Return level estimates, in line with the nature of the digital assets analyzed, indicate much greater exposure to extreme market movements for BTC and ETH than for the stablecoins.

The analysis is subsequently extended to a bivariate EVT framework to investigate extremal dependence between asset pairs. Componentwise Maxima and threshold exceedance models are estimated using both two-stage and full likelihood approaches. Significant extremal dependence is consistently identified between Bitcoin and Ethereum across these estimation methods. Evidence of dependence between USDT and USDC is weaker and not robust across specifications, while pairs combining a stablecoin with a volatile cryptocurrency show no meaningful same-tail dependence. An analysis of asymmetric tail dependence provides limited but suggestive evidence consistent with the safe haven hypothesis, at times attributed to stablecoins during periods of market stress.

Overall, the findings demonstrate that EVT provides a useful framework for quantifying tail risk in cryptocurrency markets and highlight substantial differences in both marginal tail behavior and extremal dependence between volatile cryptocurrencies and stablecoins.

Keywords: *Extreme Value Theory (EVT); Tail Risk; Cryptocurrency Markets; Generalized Extreme Value Distribution (GEV); Generalized Pareto Distribution (GPD); Multivariate Extreme Value Theory (MEVT); Safe Haven Hypothesis.*

Chapter 1

Introduction

The cryptocurrency market has experienced unprecedented growth since the establishment of Bitcoin in 2009 [53], evolving from a niche digital experiment into a global financial ecosystem with hundreds of millions of owners [12]. Initial market dynamics were almost exclusively driven by Bitcoin (BTC), originally conceived as a decentralized digital store of value. Ethereum (ETH), launched in 2015, expanded blockchain technology beyond simple payments by allowing developers to build programs, called smart contracts, directly on its network [14], a functional distinction from Bitcoin, that established it as the second pillar of the cryptocurrency market.

Much of the appeal behind this growth lies in the ability to transact directly between individuals without relying on banks or other intermediaries [11]. These cryptocurrencies, however, have mostly attracted risk-seeking investors: both BTC and ETH remain largely demand-driven assets with prices set by market supply and demand and therefore exposed to large, unpredictable swings.

For cryptocurrencies to support a broader range of portfolio allocation strategies, not only speculative positions but also more conservative ones, an instrument less exposed to this volatility was needed. Stablecoins such as Tether (USDT) and USD Coin (USDC) were introduced to fill this role by pegging their value to a reference asset, typically the US dollar, rather than letting it float freely. Their design goal is price stability rather than appreciation.

The integration of these assets has significantly diversified the market landscape and the four cryptocurrencies considered in this thesis were selected as the most representative assets of their respective categories: BTC and ETH as the two largest and most widely traded volatile cryptocurrencies; and USDT and USDC as major and widely adopted stablecoins.

This combination of structurally different assets, while sharing the same digital infrastructure but exhibiting fundamentally different price behavior, makes the cryptocurrency market a natural setting in which extreme risk can be studied. Volatile cryptocurrencies expose investors to the possibility of severe losses, while stablecoins are presumed to offer a buffer against them. Assessing this presumption requires moving beyond average market behavior to focus specifically on the tails of the return distribution, the region where standard statistical assumptions are most likely to fail.

Standard data analyses often rely on the assumption of normality, an assumption that tends to severely underestimate the frequency and severity of extreme market shocks. In the context of cryptocurrencies such events are not merely occasional outliers but structural features of the market, characterized by leptokurtosis and volatility clustering, the tendency of large price movements to be followed by further large movements. This motivates the use of statistical tools designed specifically to capture the *unusual* rather than the *usual* [21, 16]. Extreme Value

Theory (EVT) provides a rigorous asymptotic framework for modeling these rare but consequential tail events. While traditional statistics focus on the central tendency of a distribution, EVT instead characterizes the stochastic behavior of a process at unusually high or low levels, that is, extreme gains and extreme losses respectively.

EVT has already been widely applied in financial markets, particularly in the analysis of extreme co-movements across equity markets [48] and has become a standard tool for estimating tail risk measures such as return levels, Value-at-Risk [56], [25] and Expected Shortfall [34, 1] at extreme confidence levels. To satisfy the fundamental EVT requirement that observations be independent and identically distributed (i.i.d.), the literature, most notably McNeil and Frey [51], suggests a two-step filtering procedure: dynamic models such as GARCH are first applied to remove volatility clustering and EVT is then applied to the resulting standardized residuals.

This methodology has more recently been extended to cryptocurrency markets, where extreme price movements are not the exception but a recurring feature of the data [27]. Studies applying EVT to Bitcoin and other major cryptocurrencies have confirmed that their return distributions are heavy-tailed and that EVT-based risk measures provide a more reliable characterization of tail risk than measures derived from Gaussian or GARCH-only models [12].

Despite this growing body of research, an explicit gap remains in understanding the joint tail dynamics between volatile cryptocurrencies and stablecoins. Much of the existing literature focuses either on Bitcoin in isolation or on simple diversified crypto portfolios, leaving largely unexamined tail behavior of stablecoins and, in particular, their dependence structure with volatile assets during periods of extreme market distress. This is a meaningful omission: stablecoins are informally described as a safe haven [55] within the digital asset ecosystem. Yet, it is not been systematically tested using EVT whether this perceived stability holds specifically in the joint tails, as opposed to merely on average.

This thesis aims to fill that gap by investigating not only the individual tail risk of BTC, ETH, USDT and USDC, but also their bivariate dependence during periods of extreme market movements, including whether the safe haven hypothesis holds. This hypothesis states that, during extreme crashes in volatile assets such as BTC or ETH, investors rotate capital into stablecoins such as USDT or USDC. If this dynamic is present, it should appear as detectable cross-tail dependence between losses in volatile cryptocurrencies and gains in stablecoins.

Three main objectives form the backbone of the thesis. The first is outlining the marginal tail behavior of each of the four assets using univariate EVT, deriving return levels, Value-at-Risk, and Expected Shortfall as comparable risk measures across assets. The second is modeling the extremal dependence structure between all pairwise combinations of the four assets using bivariate EVT, distinguishing between same-volatility pairs, such as BTC-ETH, and mixed-volatility pairs, such as BTC-USDC. The third is testing explicitly for evidence of the safe haven hypothesis, by examining the cross-tail dependence between losses in volatile cryptocurrencies and gains in stablecoins.

The empirical strategy is conducted in two stages to achieve these objectives, following the theoretical framework developed by McNeil and Frey [51]. Standardized residuals are first obtained for each asset via GARCH-type filtering, in order to satisfy the i.i.d. assumption required by EVT. These residuals are then analyzed using both the Block Maxima and Peaks Over Threshold approaches at the univariate level, and using Componentwise Maxima and Threshold Excess models, estimated via both two-stage and full likelihood procedures, at the bivariate level followint the theory of Coles [16].

The structure of the thesis is organized as follows. Chapter 2 introduces the theoretical foundations of univariate Extreme Value Theory, and covers the Generalized Extreme Value distribution under the Block Maxima method and the Generalized Pareto Distribution under the Peaks Over Threshold method. Chapter 3 extends this framework to bivariate Extreme Value Theory, and describes how pairs of cryptocurrencies can be jointly modeled to capture tail dependence. Chapter 4 presents the data and the preliminary analysis, including the visualization of price series and the filtering procedures applied to obtain approximately independent and identically distributed returns. Chapter 5 applies the univariate methodology to the filtered return series, through both the Block Maxima approach and the Peaks Over Threshold approach. Chapter 6 applies the bivariate framework introduced in Chapter 3. Thefinal results and their limitations are discussed in Chapter 7 and the conclusions are showed in Chapter 8.

All analyses and figures presented in this thesis were performed using the `R` programming language [59] within the `RStudio` integrated development environment. GARCH filtering was carried out using the `rugarch` package [24], and both univariate and bivariate extreme value models were fitted using the `evd` package [63]. Further details on the specific functions employed are provided in the dedicated chapters.

Chapter 2

Modeling Extreme Value Theory

Extreme Value Theory (EVT) is a branch of statistics concerned with the modeling and analysis of rare, extreme events, that are occurring with very low probability. It is widely applied in fields where understanding the behavior of extreme outcomes is crucial, this includes: science, engineering, and finance [16] [21].

Extreme observations correspond to values that lie far in the tails of a distribution and depart significantly from its central tendency. In contrast to typical or *normal* events, for which a lot of data are usually available, extreme events are characterized by less observations, and extrapolating from standard models often underestimates tail risk. The important feature of EVT is that it provides quantification of the stochastic behavior of a process at unusually large or small levels [61].

In EVT, there are two principal approaches for modeling the behavior of extremes. The first is the Block Maxima method (BMM), which considers the maximum value achieved by a variable over successive, *non-overlapping* time periods, such as months or years. These maxima, referred to as Block Maxima, constitute the extreme events. For instance, in the left panel of Figure 2.1, the observations x_2, x_5, x_7, and x_{11} represent the Block Maxima for four periods, each consisting of three observations. To better align with the objectives of this study, let x_t denote the daily return of Bitcoin. Partitioning the series into non-overlapping blocks, the maximum return within each block corresponds to the largest daily movement observed, providing a sample of extreme returns for modeling tail risk. The second approach is called Peaks

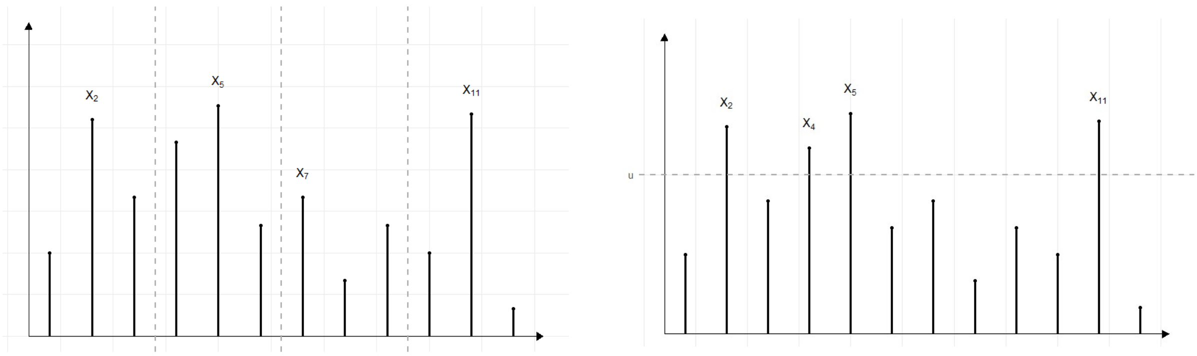

Figure 2.1: Block Maxima Method (BMM) (left) and Peaks Over Threshold (POT) (right).

Over Threshold (POT) and focuses on observations exceeding a *sufficiently high* threshold u. In the right panel of Figure 2.1, the observations x_2, x_4, x_5, and x_{11} exceed the threshold and are therefore classified as extreme events [26] [52]. In the financial context, threshold exceedance of Bitcoin returns capture unusually large market movements and provide a sample of extremes suitable for tail modeling using the POT framework.

In addition to capturing extreme gains (large positive returns), both the Block Maxima and Peaks Over Threshold approaches can be applied symmetrically to model extreme losses. By focusing on minima instead of maxima or, equivalently, by analyzing the transformed series $-X_t$, these methods allow us to study also the behavior of the left tail of the return distribution [16].

An overview of both the Block Maxima and Peaks Over Threshold methods follows in the subsequent sections.

2.1 Classical Extreme Value Theory

This first section introduces the theory that represents the basis of Extreme Value Statistics.

The focus is on the statistical behavior of the maximum of a sample,

$$M_n = \max\{X_1, \dots, X_n\},$$

where $X_1, \dots, X_n$ are independent and identically distributed (i.i.d.) random variables with common distribution function F. The variables X_i represent observations of a process recorded over time. For example, they may correspond to hourly sea-level measurements, daily average temperatures, or, as in our case, daily returns of an asset.

From a theoretical perspective, the distribution of M_n can be derived exactly. For any real number z,

$$\Pr(M_n \leq z) = \Pr(X_1 \leq z, \dots, X_n \leq z) = \{F(z)\}^n$$

For a sample of n independent observations, the probability that every value is at most z is $[F(z)]^n$.

However, this expression has limited practical use because the underlying distribution F is typically unknown. Estimating F from data can introduce errors, that increase for large n. Instead, the focus is on approximate families of distributions that describe the behavior of maxima using only extreme observations, in a way analogous to how the Central Limit Theorem [46] justifies approximating sample means by a normal distribution. Extreme Value Theory studies the limiting behavior of maxima in this sense.

To understand this behavior, we examine the limit of $\{F(z)\}^n$ as $n \to \infty$. If z lies below the upper endpoint of F, then $\{F(z)\}^n \to 0$, implying that the distribution of M_n degenerates at the upper endpoint of F. To avoid this degeneracy, we apply a linear normalization to the maxima and define

$$M_n^* = \frac{M_n - b_n}{a_n},$$

where $a_n > 0$ and b_n are sequences of constants that adjust scale and location, respectively. With appropriate choices of a_n and b_n, the normalized maxima M_n^* remain non-degenerate as n grows. Consequently, we study the limiting distribution of M_n^* rather than that of M_n directly [16].

This leads to the fundamental result of Extreme Value Theory, known as the *Extremal Types Theorem*. Introduced by *Fisher, Tippett* (1928) [23] and *Gnedenko* (1943) [29].

Theorem 1. ***Extremal Types Theorem.*** *If there exist sequences $\{a_n > 0\}$ and $\{b_n\}$ such that*

$$\Pr\left(\frac{M_n - b_n}{a_n} \leq z\right) \to G(z) \quad \text{as } n \to \infty,$$

where G is a non-degenerate distribution function, then G must belong to one of three families of distributions:

$$I\text{: } G(z) = \exp\left\{-\exp\left[-\frac{z-b}{a}\right]\right\}, \quad -\infty < z < \infty;$$

$$II\text{: } G(z) = \begin{cases} 0, & z \le b, \\ \exp\left\{-\left(\frac{z-b}{a}\right)^{-\alpha}\right\}, & z > b; \end{cases}$$

$$III\text{: } G(z) = \begin{cases} \exp\left\{-\left[-\frac{z-b}{a}\right]^{\alpha}\right\}, & z < b, \\ 1, & z \ge b. \end{cases}$$

These are known as: *Gumbel* (Type I), *Fréchet* (Type II) and *Weibull* (Type III). Each family is characterized by a location parameter b, which shifts the distribution to the right or left; a scale parameter a, which determines the spread of the distribution; and a shape parameter α, which determines the form of tail of the underlying distribution.

Rather than adopting one of the three distributions separately, a more unified approach is obtained by combining the Gumbel, Fréchet, and Weibull distributions into the Generalized Extreme Value (GEV) family, whose distribution function is

$$G(z) = \exp\left\{-\left[1+\xi\frac{z-\mu}{\sigma}\right]^{-1/\xi}\right\}, \quad \text{for } 1+\xi\frac{z-\mu}{\sigma} > 0 \tag{2.1}$$

where the parameters satisfy $-\infty < \xi < \infty$, $\sigma > 0$, and $-\infty < \mu < \infty$. The GEV distribution has three parameters: a location parameter μ, a scale parameter σ, and a shape parameter ξ.

The sign of ξ determines the type of tail behavior:

$$\begin{cases} \xi = 0 & \text{Gumbel,} \\ \xi > 0 & \text{Fréchet,} \\ \xi < 0 & \text{Weibull.} \end{cases}$$

The shapes of the probability density functions for the standard Fréchet, Weibull, and Gumbel distributions are given in Figure 2.2. We observe that the Fréchet distribution exhibits a polynomially decaying tail and is therefore well suited for heavy-tailed distributions. The exponentially decaying tails of the Gumbel distribution characterize thin-tailed distributions. Finally, the Weibull distribution arises as the asymptotic distribution for distributions with a finite upper endpoint [26]. It follows that, in applications, the three families provide distinctly different representations of Extreme Value behavior.

Bringing the three classical extreme-value families together under a single model makes statistical analysis far more straightforward. By estimating the shape parameter, ξ the data themselves reveal the appropriate tail behavior, removing the need to choose in advance among the three traditional families.

2.2 Block Maxima and Minima

To apply BMM in practice, the data are divided into non-overlapping blocks of equal size, and the maximum from each block is extracted. These Block Maxima form the sample used for inference. In practice, the normalizing constants a_n and b_n from the classical Extreme Value

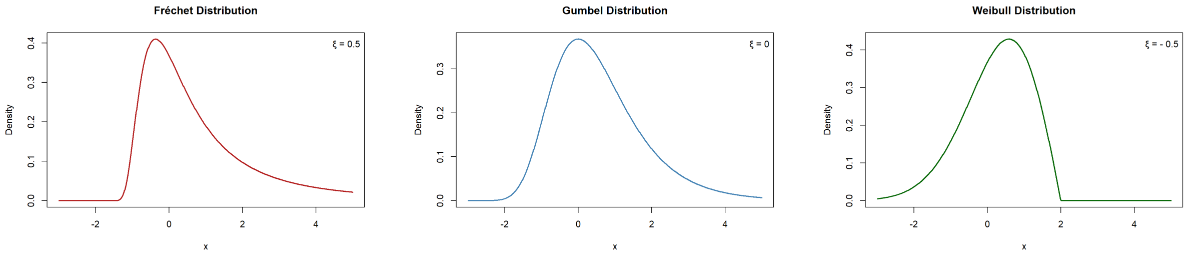

Figure 2.2: Density plots for the three Extreme Value distribution families: Fréchet (left), Gumbel (center), and Weibull (right), all with $\mu = 0$ and $\sigma = 1$.

Theorem are unknown because they depend on the underlying data distribution. Fortunately, this is not a problem: *for sufficiently large block sizes*, the Block Maxima M_n can be directly approximated by a GEV distribution.

$$\Pr\{M_n \leq z\} \approx G\left(\frac{z - b_n}{a_n}\right) = G^*(z).$$

The GEV's location and scale parameters effectively absorb the unknown constants, so explicit normalization is not needed. In other words, Block Maxima can be modeled directly with a GEV, and its parameters can be estimated from the data.

In some situations, the analysis focuses on extremely small observations rather than large ones. This applies to our case study, where we are not only interested in extreme gains but, more importantly, in extreme losses. To model this behavior, we define the Block Minima approach, which considers:

$$m_n = \min\{X_1, \ldots, X_n\}.$$

Define $Y_i = -X_i$ for $i = 1, \ldots, n$. This sign transformation implies that small values of X_i correspond to large values of Y_i. Since

$$m_n = \min\{X_1, \ldots, X_n\} \qquad \text{and} \qquad M_n^{(Y)} = \max\{Y_1, \ldots, Y_n\},$$

we obtain

$$M_n^{(Y)} = -m_n.$$

Thus, for large n, the distribution of Block Minima can be studied through the distribution of Block Maxima of the transformed variables Y_i.

If the Block Maxima of Y_i converge to a GEV distribution with parameters (μ, σ, ξ), then the Block Minima of X_i converge to a GEV distribution with parameters $(\tilde{\mu}, \tilde{\sigma}, \tilde{\xi}) = (-\mu, \sigma, -\xi)$.

Hence, for large n,

$$\Pr(M_n \leq z) \approx 1 - \exp\left\{-\left[1 - \xi\,\frac{z - \tilde{\mu}}{\sigma}\right]^{-1/\xi}\right\}, \text{ on } \quad z : 1 - \xi\,\frac{z - \tilde{\mu}}{\sigma} > 0, \text{where} \quad \tilde{\mu} = -\mu.$$

The GEV distribution for minima is obtained rewriting Equation (2.1)

$$G(z) = 1 - \exp\left\{-\left[1 - \xi\,\frac{z - \tilde{\mu}}{\sigma}\right]^{-1/\xi}\right\} \tag{2.2}$$

An alternative approach exploits the duality between maxima and minima [16]. If $Z_1, \dots, Z_m$ arise from a GEV distribution for minima with parameters $(\tilde{\mu}, \tilde{\sigma}, \tilde{\xi})$, then the transformed sample $-Z_1, \dots, -Z_m$ follows the corresponding GEV distribution for maxima. Estimating the GEV model on this transformed sample yields the parameters of the minima distribution, with the correction of the sign for $\tilde{\mu} = -\mu$.

Parameter Estimation

Several methods are available for estimating the parameters of the GEV distribution, worth noting is the Probability–Weighted Moments (PWM), widely used and well established in literature [38]. This work will, however, focus on the Maximum Likelihood approach [39][16]. Under the assumption that the Block Maxima $Z_1, \dots, Z_m$ are independent and follow a GEV distribution, the log-likelihood for the case $\xi \neq 0$ can be written explicitly as

$$\ell(\mu, \sigma, \xi) = -m \log \sigma - \left(1 + \frac{1}{\xi}\right) \sum_{i=1}^{m} \log \left[1 + \xi \left(\frac{z_i - \mu}{\sigma}\right)\right] - \sum_{i=1}^{m} \left[1 + \xi \left(\frac{z_i - \mu}{\sigma}\right)\right]^{-1/\xi},$$

which is valid only when

$$1 + \xi \left(\frac{z_i - \mu}{\sigma}\right) > 0, \quad i = 1, \dots, m$$

otherwise the likelihood is zero, since at least one observation would fall outside the support of the distribution. When the shape parameter ξ is close to zero, the general log-likelihood becomes numerically unstable, and it is therefore replaced by its Gumbel limit, which provides a well-behaved expression in the neighborhood of $\xi = 0$.

The maximization of the log-likelihood is carried out numerically, and the resulting estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi})$ are used throughout the analysis. Once the GEV parameters have been estimated, it is important to assess the uncertainty associated with the estimates through confidence intervals, computed for each of the three parameters μ, σ, and ξ.

Under standard conditions, the maximum likelihood estimators follow an approximate normal distribution. The uncertainty associated with these estimators is summarized by the variance–covariance matrix. The elements on the diagonal represent the variances of the individual parameter estimates and whose square roots correspond to their standard errors (SE). These quantities allow the construction of confidence intervals. For a 95% confidence level, the critical value of the standard normal distribution is $z_{1-\alpha/2}$, yielding the following approximate confidence interval for a generic parameter λ

$$\hat{\lambda} \pm z_{1-\alpha/2} \, \mathrm{SE}(\hat{\lambda}).$$

2.2.1 Return Levels and Financial Risk Metrics

In Extreme Value Theory, when modelling the behavior of extremes, the main objective is often the estimation of return levels, rather than the parameters themselves. Return levels quantify the magnitude of rare events over a specified time horizon [16].

Let p denote the probability that the block maximum exceeds a level z_p in a given year. The quantity z_p is the return level, and $1/p$ is the *return period*. Thus, z_p is the level expected to be exceeded, on average, once every $1/p$ periods over which the maxima are computed.

In this work, Extreme Value Theory is applied to financial markets, where related risk measures are of central importance. When the return level is computed on the left tail, when modelling losses, z_p corresponds to the Value-at-Risk (VaR) at confidence level $\alpha = 1 - p$ [51]. Value-at-Risk is the *maximum expected loss* over a target horizon such that there is a stated low probability that the actual loss will be larger [69]. In other words, VaR measures the maximum potential decline in the value of a portfolio over a given horizon at confidence level α[1]. The return level z_p is obtained by inverting the cumulative distribution function of the Generalized Extreme Value distribution in Equation (2.1), yielding

$$z_p = \begin{cases} \mu - \dfrac{\sigma}{\xi}\left[1 - \{-\log(1-p)\}^{-\xi}\right], & \text{for } \xi \neq 0, \\ \mu - \sigma \log\{-\log(1-p)\}, & \text{for } \xi = 0. \end{cases} \tag{2.3}$$

where μ, $\sigma > 0$, and ξ denote the location, scale, and shape parameters of the GEV distribution, respectively. Define

$$y_p = -\log(1-p),$$

so that Equation (2.3) can be written as

$$z_p = \begin{cases} \mu - \dfrac{\sigma}{\xi}\left(1 - y_p^{-\xi}\right), & \text{for } \xi \neq 0, \\ \mu - \sigma \log(y_p), & \text{for } \xi = 0. \end{cases}$$

When the return level is computed for gains, Equation (2.3) applies. In contrast, inverting the cumulative distribution function of the GEV for minima in Equation (2.2) yields the return level for losses:

$$z_p = \begin{cases} \tilde{\mu} + \dfrac{\sigma}{\xi}\left(1 - y_p^{-\xi}\right), & \text{for } \xi \neq 0, \\ \tilde{\mu} + \sigma \log(y_p), & \text{for } \xi = 0, \end{cases} \tag{2.4}$$

A *return level plot* is obtained by plotting z_p against y_p on a logarithmic scale. The shape of this plot provides insight into the tail behavior of the underlying distribution. In particular, when $\xi = 0$, the plot is linear, corresponding to the Gumbel case. When $\xi < 0$, the plot is convex and approaches a finite upper bound, with asymptotic limit $\mu - \sigma/\xi$ as $p \to 0$. Lastly, when $\xi > 0$, the plot is concave and exhibits no finite upper bound, indicating a heavy-tailed distribution.

By substituting the maximum likelihood estimates of the GEV parameters into both Equation (2.3) and Equation (2.4), the maximum likelihood estimate of the return level z_p associated with exceedance probability $0 < p < 1$ is obtained.

Furthermore, by the Delta Method [16], the variance of the estimated return level $\hat{z}_p$ can be approximated by

$$\mathrm{Var}(\hat{z}_p) \approx \nabla z_p(\hat{\theta})^\top \mathbf{V} \nabla z_p(\hat{\theta}),$$

where $\mathbf{V} = \mathrm{Var}(\hat{\theta})$ is the variance–covariance matrix of the maximum likelihood estimators and

$$\hat{\theta} = (\hat{\xi}, \hat{\mu}, \hat{\sigma})^\top$$

denotes the vector of estimated GEV parameters.

In Extreme Value applications, the return levels of greatest interest typically correspond to long return periods, i.e., small values of p. When $\xi < 0$, it is also possible to make inference on the upper end-point of the distribution, which corresponds to the limiting return level as $p \to 0$. In this case, the maximum likelihood estimate of the upper end-point is given by

[1] JP Morgan/Reuters. *Risk metrics.* Technical document, 1996.

$$\hat{z}_0 = \hat{\mu} - \frac{\hat{\sigma}}{\hat{\xi}} \tag{2.5}$$

while for $\xi > 0$, the upper end-point is infinite, and Equation (2.5) becomes the lower end-point estimation.

As with the parameter estimates, the uncertainty of the return level estimates is assessed through confidence intervals. Since z_p is a nonlinear function of the GEV parameters, the standard approach of adding and subtracting a multiple of the standard error is not directly applicable. Instead, a first-order Taylor expansion is used to linearize z_p and propagate the parameter uncertainty to the return level. This is known as the delta method.

The Profile Likelihood provides an alternative, likelihood-based approach to inference. It allows inference on individual parameters or functions of parameters while accounting for uncertainty in the other parameters. The profile likelihood is obtained by fixing the parameter of interest at a given value and maximizing the likelihood with respect to all remaining parameters. Repeating this procedure over a range of values for the parameter of interest produces the profile log-likelihood, which can then be used to construct approximate confidence intervals. For example, to obtain the profile likelihood for the shape parameter ξ, one fixes $\xi = \xi_0$ and maximizes the log-likelihood with respect to the remaining parameters μ and σ. Performing this step over a range of values of ξ_0 yields the profile log-likelihood for ξ, from which approximate confidence intervals can be derived [16].

This methodology can also be applied to inference on return levels. To obtain confidence intervals for a specified return level z_p, the GEV model is reparameterized so that z_p becomes one of the model parameters. Specifically, for $\xi \neq 0$,

$$\mu = z_p - \frac{\sigma}{\xi}\left[\{-\log(1-p)\}^{-\xi} - 1\right],$$

and for $\xi = 0$,

$$\mu = z_p + \sigma \log\{-\log(1-p)\},$$

so that the model is expressed in terms of the parameter vector $(z_p, \sigma, \xi)^\top$. The profile log-likelihood for z_p is then obtained by maximizing the likelihood with respect to σ and ξ, and confidence intervals follow from the resulting likelihood function.

It is important to note that all uncertainty estimates depend on the assumed model, and extrapolation to long return periods relies on assumptions that cannot be verified from the data alone. Therefore, return level estimates and their confidence intervals should be interpreted with caution, particularly for large return periods.

Expected Shortfall

To provide a complete overview of the risk measures that can be computed in EVT applications to financial analysis, we introduce the concept of Expected Shortfall (ES), also known as Conditional Value-at-Risk [60]. While Value-at-Risk identifies the threshold for rare losses, ES (CVaR) captures their severity, the magnitude, conditional on exceedance, offering a more informative picture of tail risk [1, 70]. Formally, for a given exceedance probability p, the Expected Shortfall is defined as

$$\mathrm{ES}_p = \mathbb{E}[X \mid X > z_p],$$

that is, the expected loss given that the loss exceeds the return level z_p [70], which in a financial context corresponds to the Value-at-Risk at confidence level α.

It is worth noting that in the financial risk literature, EVT and ES computation is almost exclusively developed within the Peaks Over Threshold framework, where the GPD provides closed-form expressions and a natural setting for tail exceedance modeling. The concept is nonetheless introduced here to maintain symmetry with the theoretical framework presented further on in the Threshold Models Section. For this reason the empirical and practical computation of ES in this thesis is carried out solely within the POT framework [56].

The *r* Largest Order Statistics Approach

A natural extension of the Block Maxima method is the *r* largest order statistics approach. This tries to overcome one of the biggest limitations linked to the Block Maxima framework, retaining the *r* largest observations within each block rather than only the maximum, in order to improve statistical efficiency [16]. While this generalization offers more precise parameter estimates, its advantages are largely displaced by the Peaks Over Threshold approach, which uses extreme data even more efficiently by considering all exceedances above a set threshold [52]. For this reason, the *r* largest order statistics approach is acknowledged here for completeness, but the empirical analysis in this thesis proceeds directly with the POT method.

2.3 Threshold Models

Having outlined the Block Maxima framework and its extensions, we now turn to the Peaks Over Threshold method, which forms the basis of the empirical analysis in this thesis. Rather than partitioning the data into blocks, this approach models all observations exceeding a sufficiently high threshold u, as introduced in the following. Let $X_1, X_2, \ldots$ be a sequence of independent and identically distributed random variables with marginal distribution function F. It is natural to regard as extreme events those X_i that exceed a high threshold u.

We are interested in modeling the behavior of the exceedances, defined as $y - u$.

The conditional probability that X exceeds $u + y$ given that it exceeds the threshold u is

$$\Pr(X > u + y \mid X > u) = \frac{1 - F(u+y)}{1 - F(u)}, \quad y > 0.$$

Since the parent distribution F is unknown in practice, a tractable approximation for the excess distribution is needed. This parallels the role of the GEV in the Block Maxima framework, where the parent distribution is equally unknown.

Theorem 2. *Pickands–Balkema–De Haan [6, 58]* Let $X_1, X_2, \ldots$ be a sequence of independent random variables with common distribution function F, and suppose that F satisfies the conditions of the Extremal Types Theorem, so that for large n,

$$\Pr\{M_n \leq z\} \approx G(z),$$

where G is a GEV distribution with parameters μ, $\sigma > 0$, and ξ. Then, for a sufficiently high threshold u, the conditional distribution of exceedances $y = X - u$, given $X > u$, is approximately

$$F_u(y) \approx H_{\xi,\tilde{\sigma}}(y), \qquad u \to \infty,$$

where $H_{\xi,\tilde{\sigma}}$ denotes the Generalized Pareto Distribution (GPD), defined as

$$H_{\xi,\tilde{\sigma}}(y) = \begin{cases} 1 - \left(1 + \dfrac{\xi y}{\tilde{\sigma}}\right)^{-1/\xi}, & \xi \neq 0, \\ 1 - \exp\left(-\frac{y}{\tilde{\sigma}}\right), & \xi = 0, \end{cases}$$

on $\{y : y > 0 \text{ and } (1 + \xi y/\tilde{\sigma}) > 0\}$, and where

$$\tilde{\sigma} = \sigma + \xi(u - \mu).$$

Where ξ is the shape parameter, which, as for the GEV case, determines the tail behavior of the distribution. The duality between the Generalized Extreme Value and Generalized Pareto families means that the shape parameter is dominant in determining the qualitative behavior of the Generalized Pareto Distribution, just as it is for the GEV distribution. When $\xi > 0$, the GPD has a heavy tail; when $\xi = 0$, it reduces to the exponential distribution with mean σ; and when $\xi < 0$, the distribution has a finite upper bound at $u - \tilde{\sigma}/\xi$. $\tilde{\sigma} > 0$ is the scale parameter, which controls the spread of the exceedances above the threshold u.

As the Extremal Types Theorem indicates, for a large block size n, the distribution of Block Maxima can be approximated by the Generalized Extreme Value distribution. Similarly, the Pickands-Balkema-De Haan theorem shows that for a sufficiently high threshold u, the distribution of exceedances in the Peaks Over Threshold method can be approximated by the Generalized Pareto Distribution.

Threshold selection and Parameter Estimation

Choosing an appropriate threshold is a crucial step in the Peaks Over Threshold approach. If the threshold is set too low, an excessive number of observations are classified as exceedances, increasing the variability of the parameter estimates. At high thresholds, however, the scarcity of exceedances weakens statistical reliability and may violate the asymptotic assumptions underlying the model, thereby introducing bias.

In this thesis, the threshold is selected using a combination of two complementary techniques [16].

1. The first method is based on the Mean Residual Life plot, which displays the average exceedances above a range of threshold levels u. When the Generalized Pareto Distribution provides an adequate approximation for exceedances above a threshold u_0, the mean residual life plot is expected to exhibit an approximately linear pattern in u. Confidence intervals can be added to account for sampling variability. However, in practice, mean residual life plots are often difficult to interpret and may not provide a clear indication of the optimal threshold. For this reason, in this thesis the mean residual life approach is complemented by an additional method for threshold selection.

2. The second method relies on the stability of parameter estimates across different threshold levels. The scale parameter of the GPD, $\tilde{\sigma}_u = \tilde{\sigma}_{u_0} + \xi(u - u_0)$, depends on the threshold and therefore cannot be directly used to assess parameter stability. To address this issue, a transformed scale parameter is considered, defined as $\sigma^* = \tilde{\sigma}_u - \xi u$, which is theoretically constant when the GPD model is valid.

   The Generalized Pareto Distribution is fitted to exceedances over a range of candidate thresholds, and estimates of both the shape parameter and the transformed scale parameter are obtained for each threshold. These estimates are then examined as functions of the threshold level.

   The selected threshold corresponds to the lowest value of u for which the parameter estimates remain approximately stable, allowing for sampling variability. This criterion provides a complementary and more robust basis for threshold selection.

Once a threshold u has been selected, the parameters of the Generalized Pareto Distribution can be estimated using the maximum likelihood method. Let $Y_1, \ldots, Y_k$ denote the exceedances over the threshold, where $Y_i = X_i - u$ and k is the number of exceedances.

For $\xi \neq 0$, the log-likelihood function of the GPD is given by

$$\ell(\tilde{\sigma}, \xi) = -k \log \tilde{\sigma} - \left(1 + \frac{1}{\xi}\right) \sum_{i=1}^{k} \log\left(1 + \frac{\xi y_i}{\tilde{\sigma}}\right),$$

defined for parameter values such that $(1 + \tilde{\sigma}^{-1} \xi y_i) > 0$ for all i. When $\xi = 0$, the GPD reduces to the exponential distribution and the corresponding log-likelihood has a simpler form.

Closed-form solutions for the maximum likelihood estimates are not available, so numerical optimization techniques are required. As in the GEV case, particular care must be taken when the shape parameter is close to zero and when ensuring that the estimates remain within the admissible parameter space.

As in the GEV case, standard errors and confidence intervals for the GPD parameters are obtained from standard likelihood theory, using the approximate normality of the maximum likelihood estimators. This allows the uncertainty associated with the estimated scale and shape parameters to be quantified.

2.3.1 Return Levels and Financial Risk Metrics

As in the GEV framework, extreme behavior is often interpreted in terms of return levels rather than individual parameter values. In the threshold approach, the m-observation return level x_m is defined as the level exceeded on average once every m observations.

If exceedances over the threshold u follow a GPD with parameters $\tilde{\sigma}$ and ξ, the return level is given by

$$x_m = u + \frac{\tilde{\sigma}}{\xi}\left[(m\zeta_u)^{\xi} - 1\right], \qquad \xi \neq 0,$$

where $\zeta_u = \Pr\{X > u\}$. In the special case $\xi = 0$, the expression simplifies to

$$x_m = u + \tilde{\sigma} \log(m\zeta_u).$$

In financial applications, as in the GEV case, when associated with extreme losses x_m can be interpreted as the Value-at-Risk, defined as the maximum expected loss over m observations at confidence level

$$\alpha = 1 - \frac{1}{m},$$

so that a larger return period m corresponds to a higher VaR confidence level. Return levels provide a direct and statistically grounded estimate of VaR for extreme losses, while the return level plot serves as a visual representation of risk across different probabilities or return periods.

Annualizing the return levels is straightforward: if n_y denotes the number of observations per year, the N-year return level corresponds to $m = Nn_y$. Uncertainty in GPD parameter estimates and return levels can be quantified using the delta method [65], exploiting the asymptotic normality of the maximum likelihood estimators, or the profile likelihood [57], as in the Block Maxima framework.

Thus, the POT framework offers a flexible method for estimating return levels for extreme events, complementing the BMM approach and allowing inference for return levels beyond the observed data.

Expected Shortfall

In the Peaks Over Threshold framework, the Expected Shortfall provides a natural complement to the return level and offers a more informative measure of tail risk. For a given return level x_m, interpreted as the Value-at-Risk at confidence level $\alpha = 1 - \frac{1}{m}$, the Expected Shortfall is defined as the expected loss conditional on exceeding this level.

If exceedances over the threshold u follow a Generalized Pareto Distribution with parameters $\tilde{\sigma}$ and ξ, the Expected Shortfall for a high quantile q (or equivalently for a large return period) is given by

$$\mathrm{ES}_\alpha = \frac{1}{1-\alpha}\int_\alpha^1 q_x(F)\,dx = \frac{\mathrm{VaR}_\alpha}{1-\xi} + \frac{\beta - \xi u}{1-\xi}, \qquad \xi < 1.$$

where $\beta = \tilde{\sigma} + \xi(u - u_0)$ denotes the scale parameter of the GPD re-evaluated at the threshold u, and $\xi < 1$ ensures the existence of the first moment of the tail distribution. This result holds provided that $\xi < 1$, ensuring the existence of the first moment of the tail distribution [52].

The Expected Shortfall increases more rapidly than the corresponding Value-at-Risk as the shape parameter becomes larger, reflecting the increasing thickness of the tail. In the case of heavy-tailed cryptocurrency returns, this implies that extreme losses, when they occur, are on average substantially larger than the VaR level.

Consequently, while the Value-at-Risk identifies the threshold for rare losses, the Expected Shortfall measures their expected severity, making it a more coherent and informative risk measure for Extreme Value analysis [1].

Chapter 3

Bivariate Extreme Value Theory

In the previous chapter, the univariate model was introduced to estimate the tail behavior of the four cryptocurrencies individually. In financial contexts, however, it is also of interest to analyze whether these assets are somehow related: whether a crash in one coincides with a crash in another, or whether an unexpected growth in one drives growth in the others as well [48, 28]. For this reason, the analysis is extended to multivariate Extreme Value Theory, in order to capture potential dependencies between the assets.

Following [16], the theoretical framework is developed and applied in the bivariate case. This choice is motivated by two reasons. First, the bivariate setting keeps the mathematics tractable and the models interpretable [18]. Second, parameter estimates are more reliable when the model is kept as parsimonious as possible, which is particularly important given the limited number of Block Maxima available in practice [16]. A fully parametric four-dimensional model would require significantly more parameters and stronger assumptions, moving beyond the scope of this thesis.

In practice, the bivariate framework is applied pairwise to all possible pairs of assets. With four cryptocurrencies, Bitcoin (BTC), Ethereum (ETH), Tether (USDT) and USD Coin (USDC), this yields six pairs, each analyzed independently. This pairwise approach is standard in the applied Extreme Value literature and allows the dependence structure between each pair of assets to be estimated and interpreted separately. In particular, it allows us to study how the two more volatile cryptocurrencies behave together, how the two stablecoins behave together, and whether the expected inverse relationship between stablecoins and the riskier assets is visible in the data.

The multivariate estimation procedure can follow two strategies. The first, Two-Stage Estimation, proceeds as follows:

1. Fit an appropriate univariate model (GEV via Block Maxima or GPD via Peaks-Over-Threshold) to each asset individually to characterize the marginal tail behavior.

2. Transform the original data into standard, common margins using the estimated marginal cumulative distribution functions (CDFs) to isolate and analyze the pure joint dependence structure.

The first step corresponds precisely to the univariate methodologies introduced in Chapter 2 and empirically applied in Chapter 5. The second step, which isolates the copula and the joint tail dynamics, is theoretically developed in the remainder of this chapter and applied to the data in Chapter 6. The second strategy, Full Likelihood Estimation, instead estimates all marginal and dependence parameters jointly in a single step, and is described in detail later in this chapter. Throughout the theoretical derivations, the framework is presented in terms of Maxima, extreme gains, but, by symmetry, all results apply equally to extreme losses via a standard negative sign transformation of the return series.

3.1 Copulas and Bivariate Dependence

Before introducing the formal Extreme Value framework, it is useful to connect bivariate Extreme Value Theory with Copula theory. A copula is a joint distribution function with standard uniform margins. Its primary purpose is to link individual univariate distributions into a valid joint distribution, effectively separating the pure dependence structure from the marginal behaviors [54].

To understand this mechanism, let X and Y be continuous random variables representing the daily returns of two cryptocurrencies. These assets are individually characterized by their cumulative distribution functions (CDFs), denoted as $F_X(x) = \Pr\{X \leq x\}$ and $F_Y(y) = \Pr\{Y \leq y\}$. Because each asset operates on its own scale and displays different levels of volatility, analyzing their raw returns directly would distort our understanding of their joint behavior.

To isolate how the two assets move together, we must remove these marginal differences by standardizing the data. This standardization is achieved through the Probability Integral Transform (PIT). The PIT states that if we evaluate any continuous random variable through its own CDF, the resulting values are mapped onto a continuous uniform scale between 0 and 1. Mathematically, this transformation is expressed as

$$U = F_X(X) \sim \text{Uniform}(0,1) \quad \text{and} \quad V = F_Y(Y) \sim \text{Uniform}(0,1).$$

As a result of this transform, the new variables U and V lose their original financial units and vary continuously within the interval $[0,1]$. The paired random vector (U,V) now operates entirely within the unit square $[0,1]^2$, where each component is marginally uniform. The joint distribution of this standardized uniform pair is precisely a copula. The foundational rule governing this framework is Sklar's Theorem [62].

Theorem 3. ***Sklar's Theorem.*** *Let $F(x,y)$ be a joint Cumulative Distribution Function (CDF) with continuous marginal CDFs F_X and F_Y. Then there exists a unique copula $C : [0,1]^2 \to [0,1]$ such that*

$$F(x,y) = C(F_X(x), F_Y(y)), \quad \textit{for all } x, y \in \mathbb{R}.$$

Conversely, if C is a copula and F_X, F_Y are univariate distribution functions, then $F(x,y)$ is a valid joint distribution function with marginals F_X and F_Y.

Sklar's theorem implies that any bivariate distribution can be uniquely decomposed into its individual marginal behaviors and a unique copula C, which summarizes the asset interaction. Within this domain, the parametric framework of the classical bivariate Extreme Value Logistic model corresponds to the family of Gumbel copulas [32], defined as

$$C(u,v) = \exp\left\{-\left[(-\log u)^{1/\alpha} + (-\log v)^{1/\alpha}\right]^{\alpha}\right\}, \quad \alpha \in (0,1], \tag{3.1}$$

where $u = F_X(x)$ and $v = F_Y(y)$. The parameter α explicitly modulates the degree of association: as $\alpha \to 1$, the structure simplifies to the product copula $C(u,v) = uv$, representing complete independence; on the contrary, as $\alpha \to 0$, it approaches perfect asymptotic dependence. The Gumbel copula is a member of the restricted class of Extreme Value Copulas, which are the only copula families preserved under componentwise maximization.

Specifically, we focus on the Gumbel copula, which belongs to the class of max-stable copulas. Max-stability means that if we take the maximums of our data over a certain period, the mathematical formula of their connection (the copula) stays exactly the same. This property is fundamental for multivariate Extreme Value Theory. Without max-stability, we would not have a solid mathematical foundation to predict or extrapolate the probability of extreme joint events that have never been observed in historical data.

3.2 Componentwise Maxima

Let $(X_1,Y_1),(X_2,Y_2),\ldots,(X_n,Y_n)$ be a sequence of independent and identically distributed (i.i.d.) random vectors with a joint distribution function $F(x,y)$. In our empirical setup, these vectors denote the daily returns of asset pairs, such as Bitcoin and Ethereum. Following the Block Maxima paradigm, we define the vector of Componentwise Maxima $M_n = (M_{x,n}, M_{y,n})$ where

$$M_{x,n} = \max_{i=1,\ldots,n}\{X_i\}, \quad M_{y,n} = \max_{i=1,\ldots,n}\{Y_i\}.$$

Crucially, because the maximum for each asset can occur on different days within the block, M_n represents a synthetic vector that does not necessarily coincide with any single historical observation.

Furthermore, because X and Y have different marginal distributions, modeling their raw returns together would mix up individual asset traits with their actual joint behavior. As discussed by Coles [16], standard univariate Extreme Value results apply to both components separately. Therefore, to simplify the presentation and isolate the pure dependence structure, we can transform the marginal distributions into a standard form. Following Coles [16], representations are especially simple if we assume that both variables have the standard Fréchet distribution, with the cumulative distribution function

$$F(z) = \exp(-1/z), \quad z > 0. \tag{3.2}$$

This is a special case of the GEV distribution where the parameters are fixed at $\mu = 1$, $\sigma = 1$, and $\xi = 1$. Once the margins are mapped into this common standard form, any remaining joint behavior reflects the pure, isolated structure of their tail dependence.

The limiting behavior of these standardized Componentwise Maxima is formalized by the following theorem [16].

Theorem 4. *Let $M_n^* = (M_{x,n}^*, M_{y,n}^*)$ be the rescaled vector of Componentwise Maxima, where the underlying vectors have standard unit Fréchet marginal distributions. If*

$$\Pr\{M_{x,n}^* \le x, M_{y,n}^* \le y\} \xrightarrow{d} G(x,y)$$

as $n \to \infty$, where G is a non-degenerate distribution function, then G takes the form:

$$G(x,y) = \exp\{-V(x,y)\}, \quad x > 0, y > 0,$$

where the exponent measure $V(x,y)$ is defined as:

$$V(x,y) = 2\int_0^1 \max\left(\frac{w}{x}, \frac{1-w}{y}\right) dH(w),$$

and H is a valid distribution function on the simplex $[0,1]$ satisfying the mean marginal constraint:

$$\int_0^1 w\, dH(w) = \frac{1}{2}. \tag{3.3}$$

The function $V(x,y)$, known as the Exponent Measure, fully determines the joint tail behavior. The spectral measure $H(w)$ governs how extreme mass is distributed between the two components. The constraint in equation (3.3) ensures that the marginal distributions of $G(x,y)$ remain structurally consistent with the standard unit Fréchet assumption.

Although Theorem 4 provides a complete representation of bivariate Extreme Value boundaries, the non-parametric nature of H implies that the family of valid limit distributions is infinite-dimensional, making direct empirical estimation unfeasible without further structure. To operationalize the theorem, we must adopt parametric sub-families for $V(x,y)$.

The most prominent framework is the symmetric Logistic model proposed by [32]

$$G(x,y) = \exp\left\{-\left(x^{-1/\alpha} + y^{-1/\alpha}\right)^{\alpha}\right\}, \quad x > 0, y > 0, \tag{3.4}$$

where $\alpha \in (0,1]$ represents the dependence parameter.

Substituting the standard Fréchet cumulative distribution functions $u = \exp(-1/x)$ and $v = \exp(-1/y)$ into the standard Gumbel copula (Equation 3.1)we are able to obtain the bivariate logistic Extreme Value distribution [31], reinforcing the structural unity between Gumbel copula and Logistic model.

While the logistic model is robust its main limitation, is that it is *symmetric*: it treats both components identically and cannot capture asymmetric dependence structures. Two asymmetric extensions address this limitation. The Bilogistic model [41] introduces two parameters α and β, where the difference $\alpha - \beta$ controls the degree of asymmetry. The Dirichlet (Coles–Tawn) model [17] offers an alternative asymmetric parameterization, also governed by two parameters α and β. In both cases, symmetry is recovered when $\alpha = \beta$.

Two-Stage Estimation

Following the two-stage estimation, let $z_{1,1}, \ldots, z_{1,m_1}$ and $z_{2,1}, \ldots, z_{2,m_2}$, denote the series of Block Maxima extracted from each respective asset using a monthly block size, as adopted in this analysis. By independently fitting univariate GEV distributions to these series, we obtain:

$$Z_i \sim \text{GEV}(\mu_i, \sigma_i, \xi_i), \quad i = 1,2,$$

where maximum likelihood estimation yields parameter estimates $(\hat{\mu}_i, \hat{\sigma}_i, \hat{\xi}_i)$ for each asset independently. Each marginal series is then transformed to the standard Fréchet scale via:

$$\tilde{Z}_i = \left[1 + \hat{\xi}_i \left(\frac{Z_i - \hat{\mu}_i}{\hat{\sigma}_i}\right)\right]^{1/\hat{\xi}_i}, \quad i = 1,2. \tag{3.5}$$

The transformed pairs $(\tilde{z}_{1,j}, \tilde{z}_{2,j})$ for $j \in I_{1,2}$ are then treated as independent realizations from a bivariate Extreme Value distribution with joint density:

$$g(x,y) = \left\{V_x(x,y)\,V_y(x,y) - V_{xy}(x,y)\right\} \exp\{-V(x,y)\}, \quad x > 0,\ y > 0, \tag{3.6}$$

where V_x, V_y, and V_{xy} denote the first-order partial and mixed second-order derivatives of V respectively. Given a parametric model for V with parameter vector θ (where $\theta = \alpha$ for the symmetric Logistic model), the profile log-likelihood is constructed as

$$\ell(\theta) = \sum_{j \in I_{1,2}} \log g\left(\tilde{z}_{1,j},\, \tilde{z}_{2,j};\, \theta\right). \tag{3.7}$$

Maximizing $\ell(\theta)$ yields the dependence parameter estimates along with their standard errors. A limitation of this approach is that the uncertainty in the first-stage marginal estimates is propagated into the second stage.

Full Likelihood Estimation

An alternative is to treat the marginal transformation not as a preliminary step but as part of a single joint model. Taking g to be the density of the full bivariate GEV family with arbitrary GEV margins (μ_i, σ_i, ξ_i) for $i = 1, 2$, and denoting by g_1 and g_2 the corresponding marginal densities, the log-likelihood is [16]:

$$\ell(\theta, \phi) = \sum_{j \in I_{1,2}} \log g\left(z_{1,j}, z_{2,j}\right) + \sum_{j \in I_1} \log g_1\left(z_{1,j}\right) + \sum_{j \in I_2} \log g_2\left(z_{2,j}\right), \tag{3.8}$$

where $\phi = (\mu_i, \sigma_i, \xi_i)$ for $i = 1, 2$, and $I_{1,2}$ denotes the index set of blocks in which both series have data, I_1 the blocks available only for the first series, and I_2 the blocks available only for the second. All marginal and dependence parameters are estimated simultaneously in a single optimization. As noted by [16], maximizing (3.8) corresponds to an analysis in which all available information is utilized and the potential for sharing information across variables is fully exploited. The price for this gain in statistical efficiency is an increase in computational cost.

Both strategies are applied and compared in Chapter 6. The two-stage approach preserves direct comparability with the univariate GEV results of Chapter 5, while the full likelihood approach provides more efficient joint estimates by exploiting the dependence structure across variables. Both are applied uniformly across all six asset pairs (both for losses and gains).

3.3 Bivariate Threshold Excess Model

As in the univariate case, the Block Maxima approach discards a large portion of the available data by retaining only the per-block maximum of each series. A more data-efficient alternative is to model all observations exceeding a prescribed threshold, adapting the Peaks Over Threshold framework of Section 2.2 to the bivariate setting. Recall from that section that, for a sufficiently high threshold u, the tail of an arbitrary distribution function F is approximated by the Generalized Pareto Distribution (Theorem 2). The bivariate counterpart seeks to approximate the joint distribution $F(x, y)$ on the region $x > u_x$, $y > u_y$, for suitably chosen thresholds u_x and u_y.

Let $(X_1, Y_1), \ldots, (X_n, Y_n)$ be independent realizations of a random vector (X, Y) with joint distribution function F, representing concurrent daily returns of two cryptocurrencies, such as USDT and USDC. Applying the marginal transformations

$$\tilde{X} = -\left(\log\left\{1 - \zeta_x \left[1 + \frac{\xi_x (X - u_x)}{\tilde{\sigma}_x}\right]^{-1/\xi_x}\right\}\right)^{-1} \tag{3.9}$$

$$\tilde{Y} = -\left(\log\left\{1 - \zeta_y \left[1 + \frac{\xi_y (Y - u_y)}{\tilde{\sigma}_y}\right]^{-1/\xi_y}\right\}\right)^{-1} \tag{3.10}$$

yields a transformed pair $(\tilde{X}, \tilde{Y})$ whose marginals are approximately standard Fréchet for $X > u_x$ and $Y > u_y$.

Once both marginal series have been transformed to the standard Fréchet scale via (3.9)–(3.10), the pair $(\tilde{X}, \tilde{Y})$ satisfies the conditions of Theorem 4. That theorem establishes that any non-degenerate limit distribution of rescaled Componentwise Maxima with standard Fréchet margins must take the form

$$G(x, y) = \exp\{-V(x, y)\},$$

where the exponent measure V is determined by a spectral measure H on $[0,1]$ satisfying the mean constraint

$$\int_0^1 w\, dH(w) = \tfrac{1}{2}.$$

As a consequence, the same structure applies not only to Block Maxima but also locally above the thresholds u_x and u_y: for $x > u_x$ and $y > u_y$, the joint tail of $(\tilde{X}, \tilde{Y})$ is approximated by

$$F(x,y) \approx G(x,y) = \exp\{-V(\tilde{x},\tilde{y})\}, \quad x > u_x,\ y > u_y. \tag{3.11}$$

The conditions of Theorem 4, in particular the form of V and the mean constraint on H, therefore remain in force here. The Logistic model [32] is adopted for the same reasons as in Section 3.1: it provides a parametric family for H that satisfies the mean constraint automatically, is tractable in closed form, and is flexible enough to cover all dependence levels from independence ($\alpha \to 1$) to perfect dependence ($\alpha \to 0$). As before, the main limitation of this model is its symmetry. Accordingly, the Bilogistic model [41] and the Dirichlet model [17] are also fitted, with the best-fitting specification selected via the Akaike Information Criterion (AIC) [3, 43].

A key complication of the bivariate threshold approach is that an observed pair (x_i, y_i) may exceed the threshold in only one of its components, as illustrated in Figure 3.1. For instance, USDC may exceed a given threshold while USDT does not. To handle this, the observation space is partitioned into four regions:

$$\begin{aligned}
R_{0,0} &= (-\infty, u_x) \times (-\infty, u_y),\\
R_{1,0} &= [u_x, \infty) \times (-\infty, u_y),\\
R_{0,1} &= (-\infty, u_x) \times [u_y, \infty),\\
R_{1,1} &= [u_x, \infty) \times [u_y, \infty).
\end{aligned}$$

For observations in $R_{1,1}$, both components exceed their respective thresholds and the bivariate tail model (3.11) applies directly. For observations in the remaining regions, the likelihood contribution must be *censored*, retaining only the information available from the component or components that exceed the threshold. For a point $(x,y) \in R_{1,0}$, for instance, only the x-component carries information about the joint tail, so the likelihood contribution is

$$\Pr\{X = x, Y \le u_y\} = \left.\frac{\partial F}{\partial x}\right|_{(x,u_y)}.$$

In general, the full likelihood takes the form

$$L(\theta; (x_1,y_1), \ldots, (x_n,y_n)) = \prod_{i=1}^{n} \psi(\theta; (x_i,y_i)), \tag{3.12}$$

where θ collects the parameters of the model and the contribution ψ is determined by the region in which (x_i, y_i) falls:

$$\psi(\theta;(x,y)) = \begin{cases}
\left.\dfrac{\partial^2 F}{\partial x \partial y}\right|_{(x,y)} & \text{if } (x,y) \in R_{1,1},\\[2ex]
\left.\dfrac{\partial F}{\partial x}\right|_{(x,u_y)} & \text{if } (x,y) \in R_{1,0},\\[2ex]
\left.\dfrac{\partial F}{\partial y}\right|_{(u_x,y)} & \text{if } (x,y) \in R_{0,1},\\[2ex]
F(u_x,u_y) & \text{if } (x,y) \in R_{0,0}.
\end{cases}$$

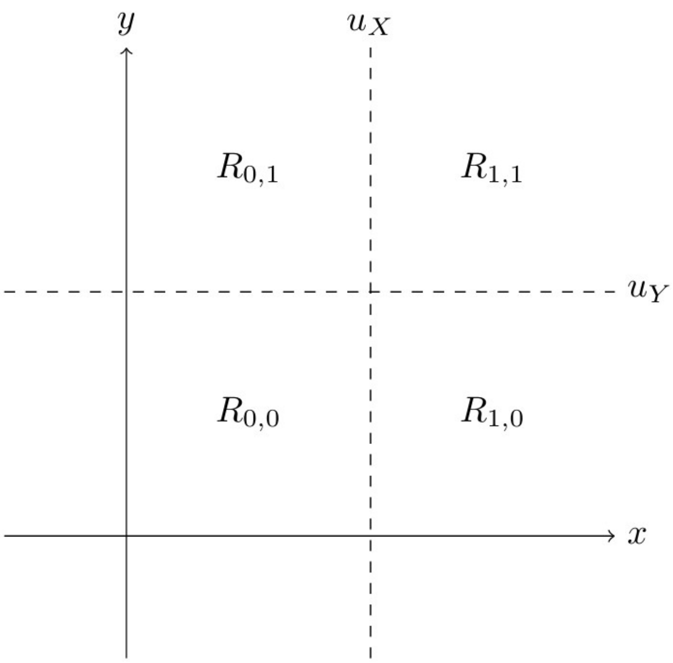

Figure 3.1: Partition of the bivariate observation space into four regions defined by the thresholds u_X (vertical dashed line) and u_Y (horizontal dashed line).

Maximizing the log-likelihood $\ell(\theta) = \log L(\theta)$ yields parameter estimates and associated standard errors in the usual way.

To implement this framework, as in the Componentwise Maxima approach, two distinct estimation strategies are deployed and compared. In the two-stage approach, the marginal parameters $(\tilde{\sigma}_x, \xi_x)$ and $(\tilde{\sigma}_y, \xi_y)$ are first estimated separately by fitting the univariate GPD to the exceedances $Y_i^x = X_i - u_x$ and $Y_i^y = Y_i - u_y$ respectively. The probability integral transformations (3.9)–(3.10) are then applied to map the data onto a standard Fréchet scale, and the censored likelihood (3.12) is subsequently maximized as a function of the dependence parameters alone ($\theta = \alpha$).

On the other hand, the full likelihood approach is fitted directly onto the original, untransformed data. In this case, the full parameter vector $\theta = (\tilde{\sigma}_x, \xi_x, \tilde{\sigma}_y, \xi_y, \alpha)$ is optimized in a single step, maximizing Equation (3.12) to estimate both the marginal GPD scaling and the asymptotic dependence structure simultaneously. This framework prevents the propagation of first-stage estimation errors and provides unified standard errors for the entire bivariated system.

3.3.1 Measures of Asymptotic Dependence: χ and $\bar{\chi}$

To visually and non-parametrically assess the tail dependence structure alongside parametric fitting, the empirical profiles of χ and $\bar{\chi}$ can be analyzed. While this technique is generally not recommended for the Block Maxima Method due to the limited number of extreme observations and consequently high sampling variance, it becomes effective and data-efficient within a threshold framework [16].

Both measures are defined on the uniform scale, obtained by transforming the marginals X and Y through their probability integral transforms $F_X(X)$ and $F_Y(Y)$, so that $u \to 1$ corresponds to moving toward the joint upper tail. This framework allows for the evaluation of the asymptotic persistence of dependence as events become increasingly extreme within the bounded domain $u \in [0,1]$. The first measure, χ, is defined as

$$\chi = \lim_{u \to 1} P\{F_X(X) > u, F_Y(Y) > u\}$$

for $0 < u < 1$

$$\chi(u) = 2 - \frac{\log P\{F_X(X) < u, F_Y(Y) < u\}}{\log u} \tag{3.13}$$

and so

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

Intuitively, $\chi(u)$ is the conditional probability that one variable exceeds a high quantile u given that the other variable also exceeds it, and χ is the limiting value of this conditional probability as the threshold is pushed to the extreme. The coefficient is bounded, $0 \leq \chi \leq 1$: a value of $\chi = 0$ indicates *asymptotic independence*, while $\chi > 0$ indicates *asymptotic dependence*, with values closer to 1 denoting a stronger tendency for extreme events to occur simultaneously in both variables.

The limitation of χ is that, whenever the variables are asymptotically independent, it collapses to zero regardless of how strong the sub-asymptotic association actually is, and it carries no information on the rate at which the dependence vanishes as $u \to 1$ [16]. To adress this, for $0 < u < 1$, $\bar{\chi}(u)$ can be defined as

$$\bar{\chi}(u) = \frac{2 \log P\{F_X(X) > u\}}{\log P(F_X(X) > u, F_Y(Y) > u)} - 1. \tag{3.14}$$

and so

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

In the regime of asymptotic dependence, χ summarizes the strength of the association. However, if the variables are asymptotically independent, χ drops to zero and fails to provide further information. In this scenario, $\bar{\chi}$ becomes the appropriate measure: it quantifies the residual dependence within the asymptotic independence regime by evaluating how fast the joint probability of extreme events vanishes compared to individual probabilities. The coefficient is bounded by $-1 \leq \bar{\chi} \leq 1$, where $\bar{\chi} > 0$ implies a positive residual association before the limit is reached, and $\bar{\chi} = 0$ signifies exact independence.

There is a direct mathematical duality between these non-parametric measures and the parametric specifications outlined in Section 3.3. Specifically, for the standard symmetric Logistic model, the tail dependence coefficient χ is explicitly tied to the dependence parameter α through the relation:

$$\chi = 2 - V(1,1) = 2 - 2^{\alpha}. \tag{3.15}$$

By virtue of (3.15), the maximum likelihood estimate of χ can be obtained directly by substituting the maximum likelihood estimate of α ($\hat{\alpha}$) into the equation[1].

[1] In the empirical analysis, this measure is computed automatically within the `fbvpot` fitting in `R`, where the output labeled *"dependence"* provides the sample estimate of χ to corroborate the parametric findings.

3.4 The Safe Haven Hypothesis: Asymmetric Tail Analysis

One of the central questions in the empirical analysis of this thesis is the existence of a flight to safety mechanism [66] within digital asset markets. Stablecoins, by their structural design, were created with the explicit purpose of offering a stable asset within the cryptocurrency market that could function as a safe investment [2]. By contrast, traditional crypto-assets such as Bitcoin and Ethereum are characterized by high-risk, high-reward dynamics [11].

This specific analysis aims to test the safe haven hypothesis [8]: namely, whether sharp declines in BTC or ETH are accompanied by significant increases in their stablecoin counterparts, USDC or USDT. In other words, the objective is to understand whether, under conditions of heightened uncertainty and market stress, investors reallocate capital away from high-risk cryptocurrencies and toward stable assets.

To empirically test this dynamic, the analysis exploits the duality between maxima and minima already introduced in the univariate analysis (Section 2.2) [16, 51], whereby the minima of a series correspond to the maxima of the sign-reversed series:

$$\min(R) = \max(-R)$$

where R represents a generic financial return. Consequently, the joint behavior of four asymmetric pairs is examined, each potentially capturing evidence of this flight to safety mechanism by pairing the losses of both BTC and ETH with the gains of USDC and USDT, respectively.

Since the objective is to capture the dependence between large losses in speculative assets (R_{crypto}) and the contemporaneous gains in stablecoins (R_{stable}), a preliminary transformation is applied to align both events onto the same positive right tail:

- **Speculative asset (Y):** $Y_t = -R_{\text{crypto},t}$ (sign reversal so that losses are analyzed as positive maxima or tail excesses).

- **Safe-haven asset (X):** $X_t = R_{\text{stable},t}$ (gains already lie naturally on the right tail).

Having defined this asymmetric sampling scheme, the subsequent modeling of the asymptotic dependence structure, together with simultaneous estimation, follows exactly the same theoretical framework described for the symmetric cases in the preceding Sections 3.3 and 3.2.

Since the full likelihood estimation approach provides a more statistically robust framework by preventing the propagation of estimation errors from the univariate marginal fits into the dependence parameters [16], the safe haven hypothesis will be tested exclusively through this estimation method.

This asymmetric methodological setup allows for isolating and testing the presence of dependence or independence in the tails. In the Block Maxima approach, it verifies whether historical maxima of losses in speculative assets tend to co-occur with historical maxima of gains in stablecoins within the same temporal horizons. Symmetrically, within the bivariate threshold excess model, it evaluates the joint probability of exceeding critical thresholds u_x and u_y, allowing for the comprehensive analysis of the joint tail behavior under the exact same conditional framework.

Chapter 4

Data and Preliminary Analysis

4.1 Data Description

This section describes the datasets used in the empirical analysis. For the scope of the analysis four *OHLC* (Open, High, Low, Close) cryptocurrencies time series were collected: Bitcoin (BTC), Ethereum (ETH), Tether (USDT) and USD Coin (USDC), with all prices expressed in U.S. dollars (USD). For this study, only daily closing prices are used. Figure 4.1 displays the evolution of these series.

All datasets were downloaded in CSV format from publicly available sources, Bitcoin data were collected from *Investing.com*[1], starting from the earliest available records [27]. Ethereum, Tether, and USD Coin data were downloaded from *CoinMarketCap.com*[2].

Bitcoin (BTC)

Bitcoin [53] is the main asset in this analysis. The dataset includes daily OHLC prices, from 20 July 2010 to 18 January 2026, covering the entire history of the market, from its early low-liquidity stage to its current global adoption.

Because Bitcoin has the largest market capitalization and serves as the primary gateway to the crypto ecosystem, it is used as the benchmark for measuring volatility and shock transmission [44].

Ethereum (ETH)

Ethereum [14] is included as the second major cryptoasset in the analysis, using daily data from 7 August 2015 to 18 January 2026. While Bitcoin is mainly held as a store of value, Ethereum functions both as a cryptocurrency and as a payment token within its own network, thereby reflecting broader market activity [15]. For the purposes of this thesis, however, Ethereum is considered solely in its role as a cryptocurrency. Including Ethereum enables us to capture wider crypto-market dynamics beyond Bitcoin, allowing us to distinguish Bitcoin-specific behavior from market-wide movements.

[1] https://www.investing.com/
[2] https://coinmarketcap.com/

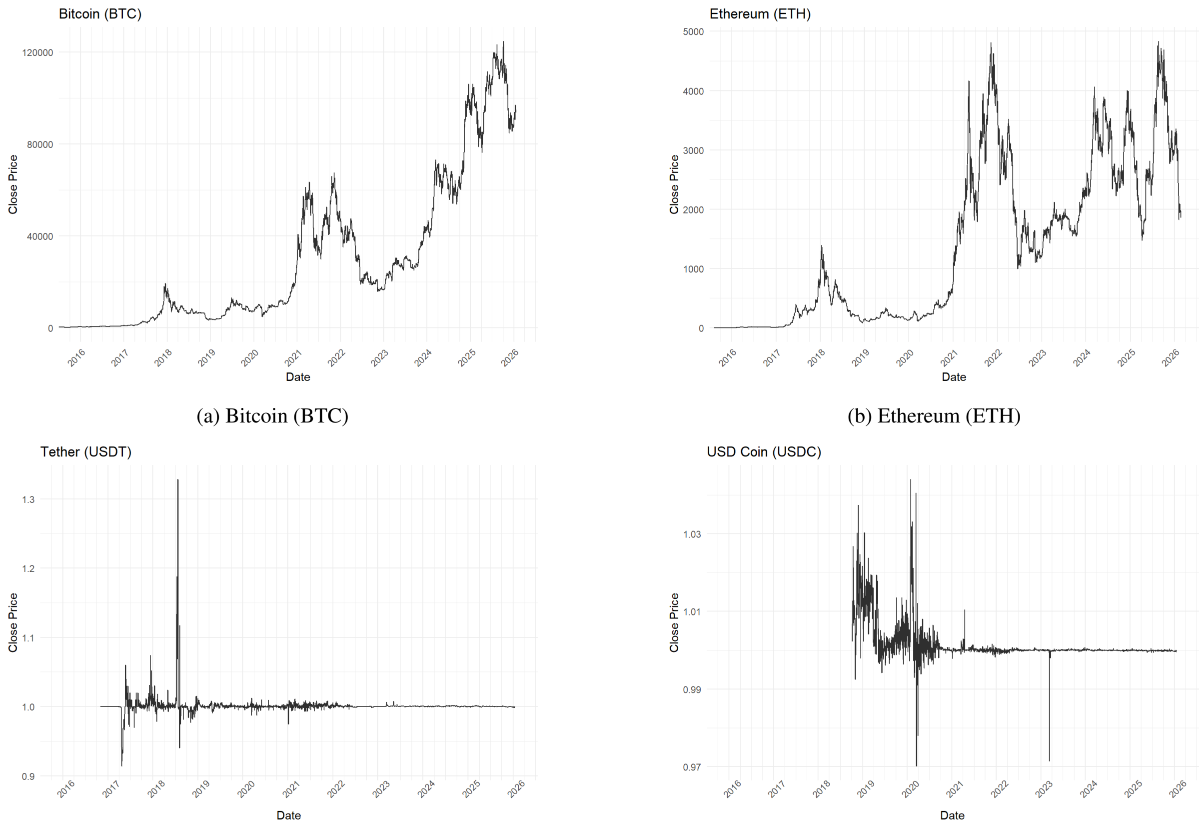

(a) Bitcoin (BTC)

(b) Ethereum (ETH)

(c) Tether (USDT)

(d) USD Coin (USDC)

Figure 4.1: Daily closing prices of the four cryptocurrency datasets used in the analysis. All series are shown over the period January 2016 – January 2026 to facilitate visual comparison across assets.

Stablecoins: USDT and USDC

To compare volatile assets with low-volatility instruments, two stablecoins are included:

- **Tether (USDT):** 2 November 2016 – 18 January 2026
- **USD Coin (USDC):** 8 October 2018 – 18 January 2026

Stablecoins, like the name suggests, are designed to maintain a price close to 1 USD, commonly referred to as the *peg* [30, 40], which leads to time series that tend to fluctuate within a narrow band around this value. The peg is supported by three main mechanisms:

- *Convertibility:* coins can be created or redeemed at a fixed price of approximately 1 USD;
- *Reserves:* issuers hold assets intended to support the fixed redemption value;
- *Arbitrage:* traders exploit small deviations from 1 USD, rapidly pushing the market price back toward the peg.

Data of both stablecoins are collected from the early stages of their respective cryptocurrencies. As visible in Figure 4.2, the behavior of stablecoins in their initial period is far from stable. In the first years after launch, market mechanisms made it difficult for the peg mechanism to function properly and the trust necessary to anchor the price to the reference value had yet to consolidate. For USDT we observe the largest upward deviation at approximately $+0.3$ above

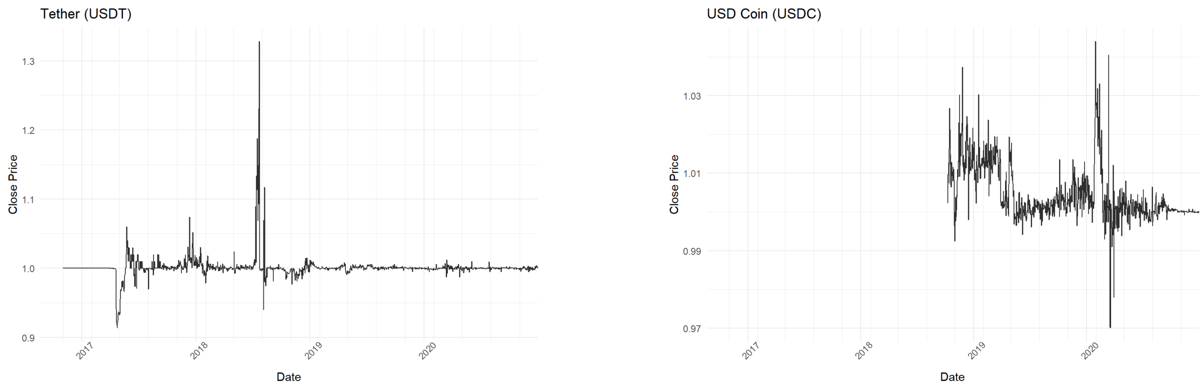

(a) Tether (USDT)

(b) USD Coin (USDC)

Figure 4.2: Daily closing prices of the two stablecoins, truncated to the 2016–2020 period to highlight their early-stage instability before converging toward a more stable peg.

the target value of 1 USD. For USDC, deviations are also measured relative to the 1:1 peg with the US dollar, but remain considerably smaller, with a maximum positive deviation of approximately $+0.035$ and a maximum negative deviation of approximately -0.03.

After these early-stage fluctuations, stablecoin prices exhibit very limited variability, with only some significant movements from the peg. This is the feature that is most relevant to this thesis, and the reason behind wanting to add stablecoins to the analysis. Stablecoins often attract demand during periods of high crypto-market volatility [7, 50]. When uncertainty rises in the market, investors move capital from volatile assets, like Bitcoin and Ethereum, into stablecoins, which act as a low-volatility liquidity option within the crypto ecosystem. We already referred to this mechanism as the flight to safety mechanism [66] or the safe haven hypothesis [8].

4.2 Distributional Properties

For the purposes of this analysis, daily closing prices are transformed into daily log returns [64], defined as

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \tag{4.1}$$

where P_t denotes the closing price at day t (Figure 4.3). This transformation is standard in financial time series analysis as it ensures approximate stationarity and yields a scale-free measure of price changes comparable across assets. All subsequent analysis are conducted on the log return series.

When computing the log returns, a problem came up with USDT. In its early stages, from the end of 2016 to approximately March 2017, the stablecoin's price was fixed at 1 and did not change (visible at the beginning of the left panel of Figure 4.2). Computing log returns as in Formula 4.1 therefore yields all zeros over this period, since $\ln(P_t/P_{t-1}) = \ln(1) = 0$ whenever consecutive prices are identical. Given that the preparatory analysis performed in the following paragraphs to ensure that returns are i.i.d. can be distorted by a long sequence of zeros clustered at the beginning of the series, these observations are dropped, and the USDT dataset used throughout the paper starts on 26 March 2017.

After this truncation, a small number of zero returns remain in the series, less than 40 observations out of 3,221 in total. These are again caused by periods in which the price did not change from one day to the next. Unlike the initial block of zeros, however, these residual observations are spread across the full sample and do not concentrate in any particular region of the distribution, so they are unlikely to distort the subsequent estimation and are therefore left in the dataset.

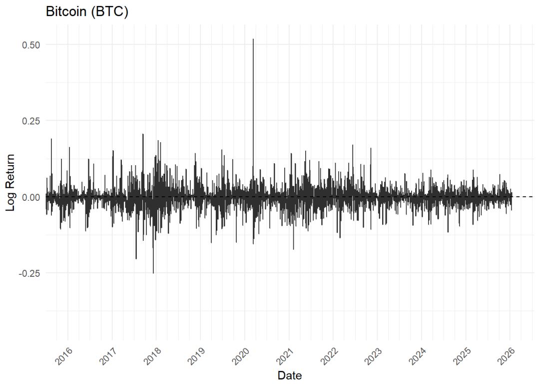

(a) Bitcoin (BTC)

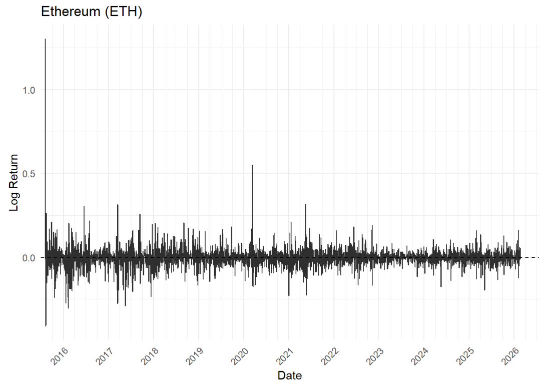

(b) Ethereum (ETH)

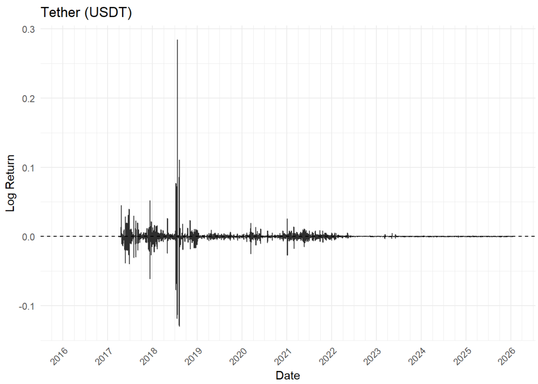

(c) Tether (USDT)

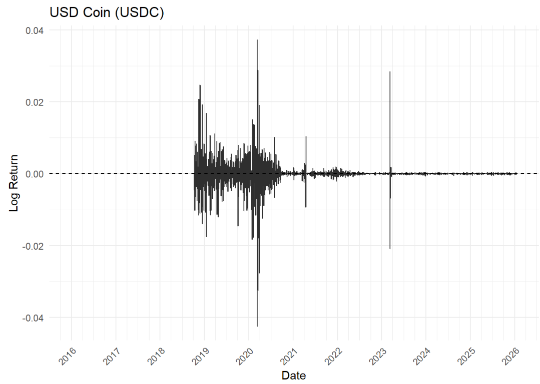

(d) USD Coin (USDC)

Figure 4.3: Daily log returns of the four cryptocurrency datasets used in the analysis. All series are shown over the period January 2016 – January 2026 to facilitate visual comparison across assets.

Before proceeding to the Extreme Value analysis, a preliminary exploratory analysis is carried out to assess whether the return series exhibit heavy-tailed behavior, a key prerequisite for the application of EVT, which is specifically designed to model distributions whose tails decay more slowly than those of the normal distribution [21, 16].

For each asset, a histogram of the empirical return distribution is plotted alongside the fitted normal density curve, and a normal QQ plot is produced (Figure 4.4). Leptokurtic behavior manifests in the histograms as a sharper and taller central peak than the normal curve predicts, with observable mass in the tails beyond what the normal distribution would suggest. In the QQ plots, heavy tails appear as a characteristic S-shaped deviation from the reference line, where both the lower left and upper right tails bend away from normality. All four assets display this pattern clearly.

To Quantify the tail heaviness, the Kurtosis measure is computed for each series. Distributions can be classified into three types based on this measure: leptokurtic (kurtosis > 3), mesokurtic (kurtosis ≈ 3, as in the normal distribution), and platykurtic (kurtosis < 3) [19]. The results are reported in Table 4.1.

Table 4.1: Kurtosis of the log returns of the four cryptocurrencies under analysis.

Asset	Kurtosis
Bitcoin	20.03
Ethereum	82.19
USDT	377.69
USDC	55.62

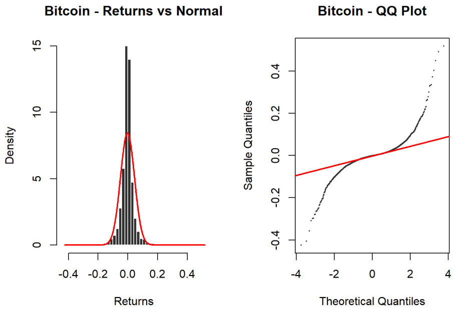

(a) Bitcoin (BTC)

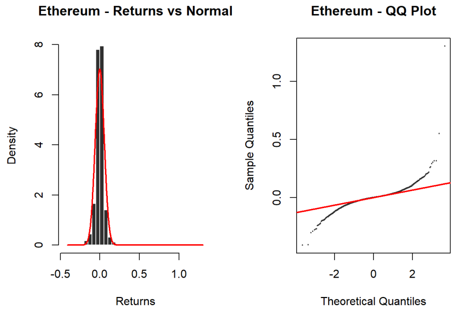

(b) Ethereum (ETH)

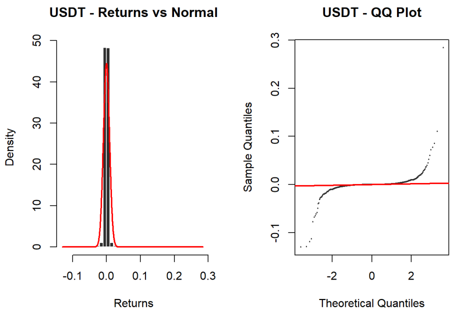

(c) Tether (USDT)

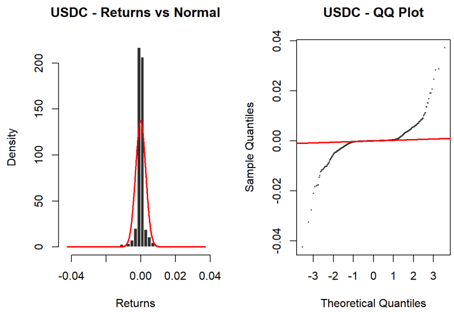

(d) USD Coin (USDC)

Figure 4.4: Histogram and QQ plot of log returns for the four cryptocurrency series. The red curve represents the fitted normal distribution.

All four assets display kurtosis values far exceeding 3, confirming strongly leptokurtic behavior. The values range from 20 for Bitcoin to nearly 378 for USDT. These results confirm that the normal distribution substantially underestimates the probability of extreme events across all four series, providing clear empirical motivation for the application of Extreme Value Theory [21, 16, 51].

4.3 Stationarity and Independence Diagnostics

Both the Block Maxima and the Peaks Over Threshold approaches require i.i.d. series, so independent and identically distributed observations. To ensure this, some preprocessing steps were carried out in order to correctly apply the theory explained in Chapter 2. While the previous section confirmed that the return series exhibit heavy-tailed behavior suitable for EVT, log returns are not automatically independent and identically distributed. Financial return series often exhibit dependence structures, particularly in their volatility, even when the returns

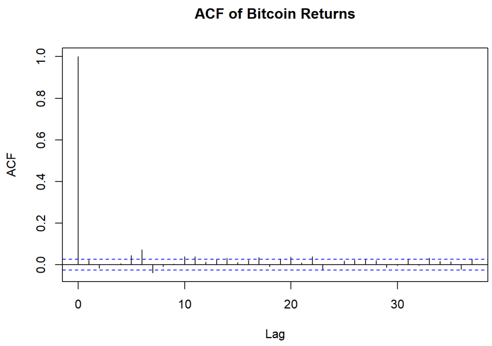

(a) Bitcoin (BTC)

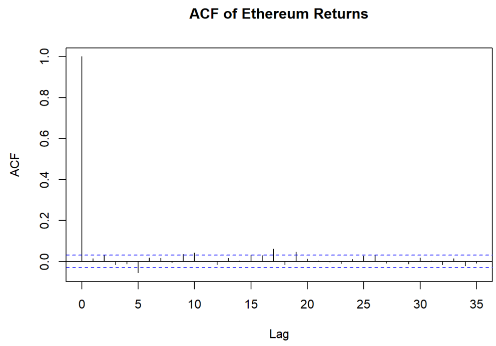

(b) Ethereum (ETH)

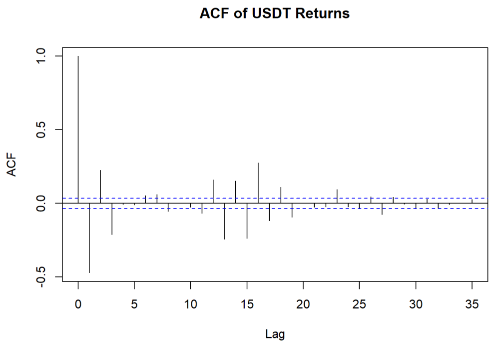

(c) Tether (USDT)

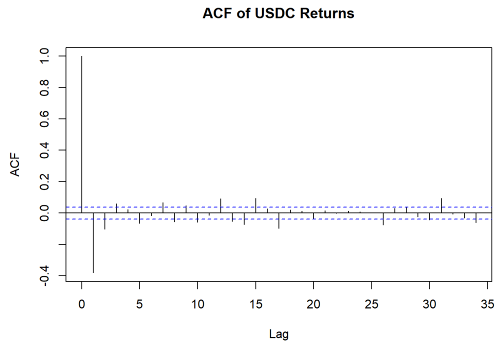

(d) USD Coin (USDC)

Figure 4.5: ACF of the returns of the four cryptocurrencies used in the analysis.

themselves appear uncorrelated. Therefore, additional diagnostics are required to assess the dependence properties of the data.

Initially the autocorrelation function (ACF) of the returns was examined (Figure 4.5).

This helps us understand whether there is correlation in the mean.

For both Bitcoin and Ethereum, the vast majority of autocorrelations across all lags fall well within the confidence bands, confirming that serial correlation in the mean is negligible. With only a marginal spike at lag 6 that significantly reaches the upper bound for Bitcoin, and three small but statistically significant spikes (at lags 10, 17 and 19) for Ethereum. The ACF structure of both series is therefore indistinguishable from white noise, and serial correlation in the mean is considered absent.

For the stablecoins USDT and USDC, the picture differs notably. USDT displays an irregular ACF structure with prominent spikes around lags 1–3 and again near lags 12–18, with several additional lags exceeding the confidence bounds. USDC shows a negative spike at lags 1–2, alongside further exceedances scattered across the ACF. Both stablecoins thus exhibit meaningful serial correlation in the mean, standing in clear contrast to the white noise behavior observed for Bitcoin and Ethereum.

Beyond autocorrelation in the mean, the ACF of the squared returns was examined (Figure 4.6) to analyze possible autocorrelation in the variance. For Bitcoin, significant autocorrelations are visible at the first several lags with a gradual decay, providing clear evidence of volatility clustering and conditional heteroskedasticity [4]. Ethereum displays a considerably weaker pattern, with only lags 3 and 5 marginally exceeding the bounds, suggesting mild but present volatility persistence. For USDT, pronounced spikes appear at early lags and again between lags 12 and 17, while USDC shows exceedances throughout the entire plot. Despite their

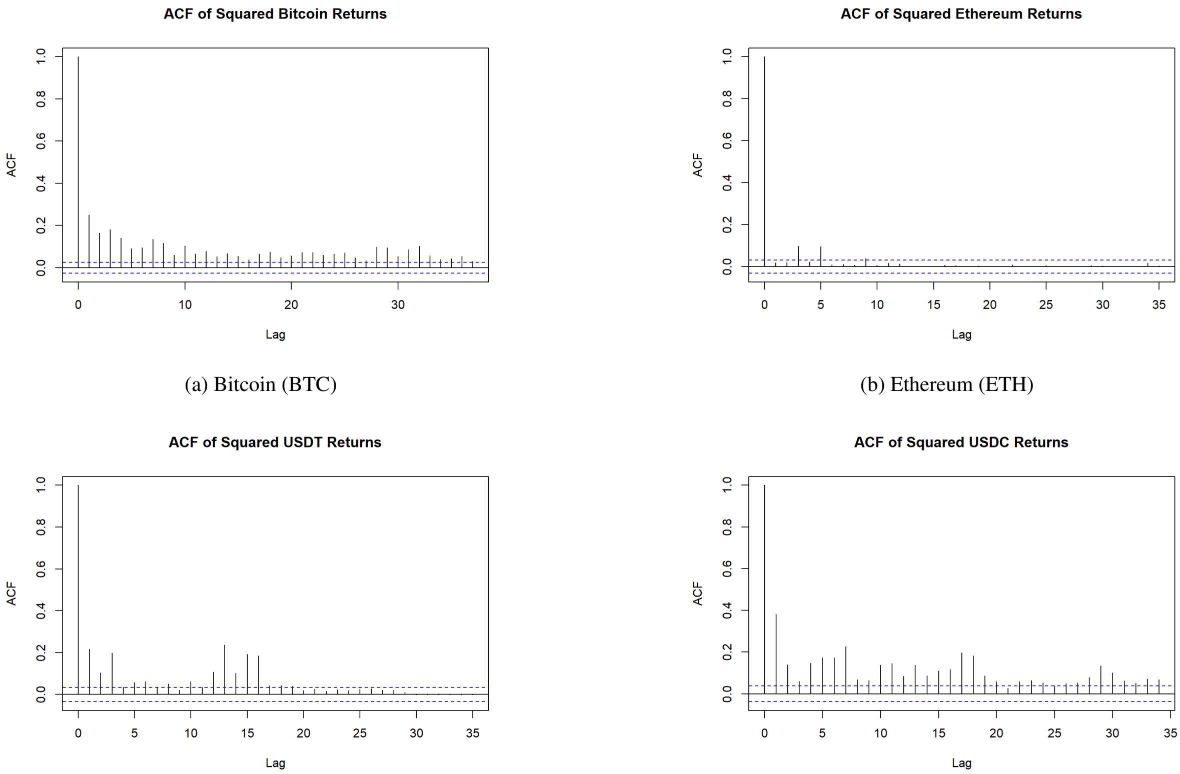

(a) Bitcoin (BTC)

(b) Ethereum (ETH)

(c) Tether (USDT)

(d) USD Coin (USDC)

Figure 4.6: ACF of the squared returns of the four cryptocurrencies used in the analysis.

price-stabilizing design, both stablecoins therefore exhibit substantial volatility clustering, with dynamics markedly different from those of Bitcoin and Ethereum.

In order to test the independence condition of the log returns, we perform two statistical tests: the Ljung-Box test and the Runs test.

The first is the Ljung-Box test [47], which evaluates the null hypothesis H_0 of absence of serial correlation up to a given lag m, against the alternative hypothesis H_1 that at least one autocorrelation coefficient is significantly different from zero. Table 4.2 reports the p-values obtained at 10 and 20 lags. The null hypothesis is rejected, since we observe a small p-value ($p < 0.001$), for all four analyzed assets. While this result is consistent with the presence of linear dependencies observed in the ACF plots of the stablecoins (USDT and USDC), the rejection for Bitcoin and Ethereum requires a different interpretation. For these two assets, the visual analysis of the ACF of the residuals (higher panels of Figure 4.5) suggested a behavior close to white noise. In light of this discrepancy between the formal test result and the visual evidence, and in order to avoid overfitting the conditional mean equation with dynamics that could be spurious, we treat the mean as uncorrelated for Bitcoin and Ethereum.

The second test applied is the non-parametric Runs test proposed by Wald and Wolfowitz [68]. This methodology verifies the hypothesis of independence by analyzing the sequence of positive and negative returns. A run is defined as a consecutive sequence of returns with the same sign. Under the null hypothesis H_0, the number of runs follows a random distribution, whereas the alternative hypothesis H_1 implies non-random behavior due to excessively long runs, indicating persistence, or excessively short runs, indicating frequent alternation. As shown in Table 4.2, the Runs test rejects the null hypothesis for all assets, indicating that the directional sequence of returns deviates from pure randomness.

Table 4.2: P-values of independence diagnostic tests on log returns.

Asset	Ljung–Box (10)	Ljung–Box (20)	Runs Test
Bitcoin	0.0000	0.0000	0.0072
Ethereum	0.0004	0.0000	0.0003
USDT	0.0000	0.0000	0.0000
USDC	0.0000	0.0000	0.0000

Taken together, the visual inspection of the ACF plots and the formal statistical tests provide strong and consistent evidence that the raw log return series cannot be treated as i.i.d. The presence of serial dependence and volatility clustering violates both the independence and the identical distribution assumptions required for the application of Extreme Value Theory.

These violations are not merely a technical inconvenience, they reflect a fundamental characteristic of financial time series that must be explicitly accounted for before any Extreme Value analysis can be conducted. Simply applying the Block Maxima or Peaks Over Threshold methods to raw returns would risk producing unreliable estimates, as the theoretical guarantees of EVT rely on the i.i.d. assumption [16]. A natural and widely adopted solution in the financial econometrics literature is to first filter the return series through a GARCH-type model [51, 33], which explicitly models the time-varying volatility. This filtering procedure is described in detail in the following section.

4.3.1 GARCH Filtering

The preliminary analysis presented in the previous section provides clear evidence that the log return series as it is cannot be treated as i.i.d. In particular, the ACF of returns revealed significant autocorrelation in the mean for the two stablecoins and the ACF of the squared returns revealed volatility clustering, autocorrelation in the variance, across all four assets: large price movements tend to be followed by further large movements, and calm periods tend to be followed by further calm periods.

This patterns violate the independence and identically distributed assumption required for the direct application of Extreme Value Theory.

To address this, a GARCH-type filtering procedure is applied before the EVT analysis, using the `rugarch` package [24] in `R` [59], following the methodological framework presented in the course materials by Embrechts et al. [20], which are themselves grounded in the foundational texts of the field [52, 37]. These references provide the practical foundation for both the GARCH filtering procedure and the subsequent univariate EVT analysis conducted in this work.

ARCH and GARCH Models Theoretical Framework

To understand the GARCH model, it is useful to introduce first its predecessor, the Autoregressive Conditional Heteroskedasticity (ARCH) model [22]. The key intuition behind ARCH is that this time-varying variance is not random: it depends on how large the most recent shocks have been. A large unexpected return today, whether positive or negative, tends to be followed by further periods of elevated volatility.

In other words, ARCH captures the idea that turbulent times in financial markets tend to cluster together.

While ARCH was an important conceptual step forward, it has a practical limitation: capturing long periods of volatility persistence requires including a large number of past shocks in the model, making it parameter-heavy and difficult to estimate reliably. Bollerslev [9] addressed this by proposing the Generalized ARCH, or GARCH model. The key improvement is that GARCH models volatility is depending not only on past shocks, but also on its own past values. Intuitively, if volatility was already high yesterday, it is likely to remain elevated today, regardless of whether a new shock has occurred. This means that a single parsimonious specification, the GARCH(1,1) using just one lag of each component, is sufficient to capture the long memory in volatility that characterizes most financial time series [64].

Formally, the GARCH(1,1) model is specified as

$$\begin{aligned} r_t = c + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t, \qquad z_t \overset{i.i.d.}{\sim} N(0,1), \\ \sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \end{aligned}$$

Where r_t is the return at time t, and it is decomposed into a constant mean c and a shock term ε_t. The shock is itself the product of two components: a time-varying standard deviation σ_t, which captures the current level of volatility, and a standardized innovation z_t, which is an i.i.d. draw from a standard normal distribution and carries the purely random, unpredictable component of the return. Crucially, while z_t is i.i.d., the shock ε_t is not, because its variance σ_t^2 changes over time according to the second equation. There, $\omega > 0$ represents the long-run average level of variance; $\alpha \geq 0$ measures how strongly a large shock yesterday increases volatility today; and $\beta \geq 0$ measures how much of yesterday's volatility level carries forward into today. The constraint $\alpha + \beta < 1$ ensures covariance stationarity, meaning that volatility fluctuates around a finite long-run level rather than growing without bound.

Two-Step Filtering Procedure

The methodology adopted throughout this thesis follows a two-step filtering framework that is standard in the financial econometrics literature. In the first step, a GARCH(1,1) model is fitted to each return series to remove serial dependence and render the residuals approximately i.i.d. In the second step, detailed in the following chapter, Extreme Value Theory is then applied to these standardized residuals, which by construction are free from the dependence structures that would otherwise violate the underlying EVT assumptions [51].

Given the strong leptokurtic behavior documented in the previous section, relying on the classical normally distributed innovations would likely underestimate the probability of extreme events. For this reason, Student-*t* innovations are adopted instead, since the heavier tails of the t-distribution offer a more realistic characterization of return shocks.

Innovators are then expressed as $z_t \overset{i.i.d.}{\sim} t_\nu(0,1)$, where z_t are still i.i.d. variables with zero mean and variance equal to one, and where ν is the degrees of freedom parameter. ν is not fixed *a priori* but is estimated jointly with the GARCH parameters via maximum likelihood estimation, allowing the data themselves to determine the appropriate level of tail heaviness. To validate, also empirically, this choice, both models, with both Normal and Student-*t* specification were estimated and compared using the Akaike Information Criterion (AIC) [3, 43]. Across all four assets, the GARCH model with Student-*t* innovations consistently yielded a lower AIC than the Normal counterpart, confirming that the Student-*t* specification provides a statistically superior fit to the data [13]. The Normal distribution model was therefore discarded in favor of the *Student-t* for all subsequent analysis. Once the model has been estimated, the standardized residuals are extracted as

$$\hat{z}_t = \frac{\varepsilon_t}{\hat{\sigma}_t} = \frac{r_t - \hat{c}}{\hat{\sigma}_t}, \tag{4.2}$$

where $\hat{c}$ is the estimated constant mean and $\hat{\sigma}_t$ is the estimated conditional standard deviation at time t. By dividing each return by its own estimated volatility, this transformation removes the volatility clustering from the series. This two-step approach has become a standard reference in the literature on extreme risk measurement [21, 52] and has been widely applied to cryptocurrency studies [42, 27, 12].

For Bitcoin and Ethereum, the GARCH(1,1) specification with Student-t innovations performs well. After filtering, the ACF of the squared standardized residuals shows no significant autocorrelation confirming that the volatility clustering has been successfully removed, as shown in Figure 4.7.

The standardized residuals of these two assets are therefore considered approximately i.i.d. and are used directly as input for the subsequent EVT analysis.

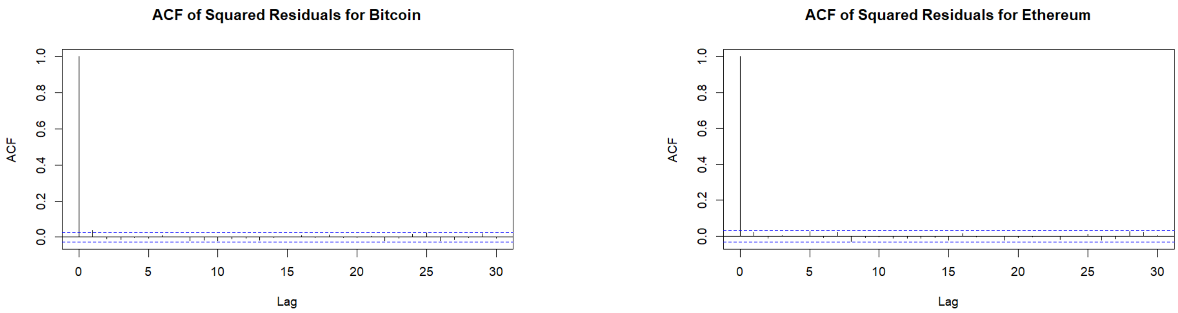

(a) Bitcoin (BTC)

(b) Ethereum (ETH)

Figure 4.7: ACF of standardized squared residuals after GARCH(1,1) filtering for Bitcoin and Ethereum.

For the two stablecoins, the situation is more complex. As already documented in the ACF analysis of raw returns, USDT and USDC exhibit significant dependence not only in their squared returns, but also in the returns themselves. This is a direct consequence of the peg stabilization mechanism: when the price deviates from 1 USD, the three mechanisms immediately push the price back toward the peg [50, 30]. This creates a predictable mean-reverting pattern in returns, a positive deviation today is systematically followed by a negative correction tomorrow, that a standard GARCH model with a constant mean cannot capture, because GARCH only models the variance and leaves the mean equation unchanged.

To address this, an ARMA(1,1) component is added for both USDT and USDC. An ARMA model captures linear dependence in the level of the returns themselves [10]. In simple terms, the ARMA component models and removes the predictable mean-reverting behavior induced by the peg mechanism, while the GARCH component, still with Student-t innovations, continues to handle the volatility clustering. The two components therefore serve complementary roles: ARMA cleans the mean, GARCH cleans the variance. The formal model of the ARMA(1,1) is:

$$\begin{aligned} r_t &= c + \phi\, r_{t-1} + \theta\, \varepsilon_{t-1} + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t, \qquad z_t \overset{i.i.d.}{\sim} t_\nu(0,1), \\ \sigma_t^2 &= \alpha_0 + \alpha_1\, \varepsilon_{t-1}^2 + \beta\, \sigma_{t-1}^2, \end{aligned}$$

Consequently, for USDT and USDC, the standardization of the residuals must account for this time-varying conditional mean. The standardized residuals are therefore obtained as:

$$\hat{z}_t = \frac{\varepsilon_t}{\hat{\sigma}_t} = \frac{r_t - \hat{E}_{t-1}[r_t]}{\hat{\sigma}_t}, \tag{4.3}$$

Where $\hat{E}_{t-1}[r_t] = \hat{c} + \hat{\phi}\, r_{t-1} + \hat{\theta}\, \varepsilon_{t-1}$ that corresponds to the estimated conditional mean of the return at time t, computed using all information available up to time $t-1$, and $\hat{\sigma}_t$ is the conditional standard deviation estimated by the GARCH process.

For completeness, given the residual uncertainty regarding the correlation structure in the mean of BTC and ETH discussed earlier, an ARMA(1,1)–GARCH(1,1) specification was additionally estimated for both these assets, in order to verify whether it provided a better fit than the simple GARCH(1,1). The two specifications were compared using the AIC, consistent with the model selection criterion adopted throughout this thesis. For both assets, the GARCH(1,1) model yielded a lower AIC, confirming the choice made at the beginning of this section to retain a constant-mean GARCH(1,1) specification for BTC and ETH.

As a methodological detail, the raw returns of both stablecoins were multiplied by 100 prior to estimation. This solution was adopted because stablecoin price fluctuations are extremely small. In their original decimal form, these tiny numbers caused numerical instabilities and optimization errors during the GARCH estimation in R [59]. Rescaling the data to percentages resolves this issue without affecting the standardized residuals or any subsequent statistical inference. The only consequence of this adjustment is that the calculated return levels must be divided by 100 afterward to restore their original real-world magnitude.

After applying the ARMA(1,1)–GARCH(1,1) specification, the squared residuals of both stablecoins show no significant autocorrelation, confirming that the volatility clustering has been successfully removed, as showed in Figure 4.8. Table 4.3 displays the p-values for the Ljung-Box Test on the standardized residuals of the stablecoins. As shown in the table, for lag 5, 10 and 20 we can not reject (p-value > 0.01) the null hypothesis of no serial autocorrelation in the mean.

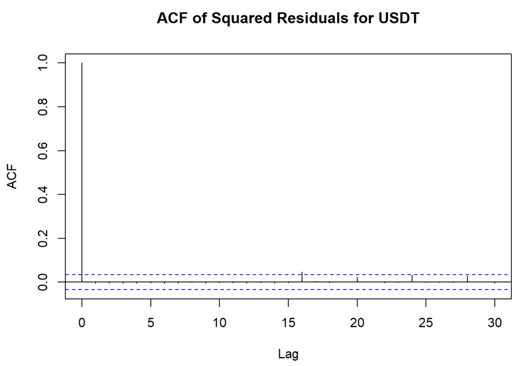

(a) Tether (USDT)

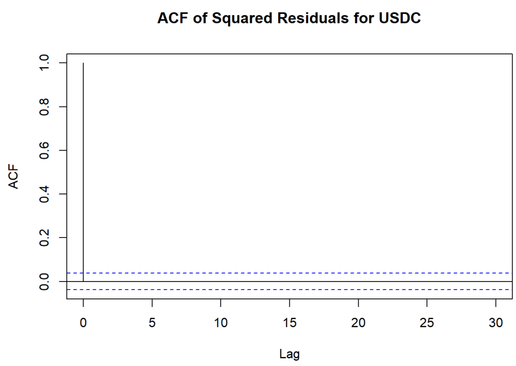

(b) USD Coin (USDC)

Figure 4.8: ACF of standardized squared residuals after ARMA-GARCH filtering for USD Coin and Tether.

Table 4.3: Ljung–Box test on standardized residuals of USDC Coin and Tether.

Asset	Lag	p-value
USDT	5	0.2640
	10	0.4954
	20	0.5500
USDC	5	0.9888
	10	0.9995
	20	0.9999

Overall, the GARCH(1,1), and ARMA(1,1)–GARCH(1,1), filtering procedure has proven effective across all four assets, producing standardized residuals that are approximately i.i.d. These filtered residuals satisfy, at least approximately, the independence assumptions required by Extreme Value Theory and are therefore used as the input for the Block Maxima and Peaks Over Threshold analyses carried out in the following chapter.

Finally, it is worth acknowledging that a range of more advanced GARCH models exist beyond the basic one adopted here. More sophisticated variants, capable of capturing asymmetric effects, non-linear dynamics, and complex conditional distributions, might have more closely reflected unique cryptocurrency behaviors [42]. However, this analysis intentionally focuses on foundational models to maintain a clear and interpretable baseline, and this remains a recognized limitation of the work.

Chapter 5

Univariate Extreme Value Theory Analysis

The univariate Extreme Value analysis was conducted in `R` [59], primarily using the `evd` (Extreme Value Distributions) package [63], which provides functions for fitting Generalized Extreme Value models and Generalized Pareto Distribution models, computing profile likelihoods, and extracting diagnostic plots.

5.1 Block Maxima and Minima Approach

Following Section 2.2, the Block Maxima Method is applied to the GARCH-filtered residuals of all four assets, extracting Block Maxima and Minima to capture extreme losses and gains respectively. A block size of 30 calendar days is adopted, reflecting the fact that cryptocurrencies trade continuously and, unlike traditional financial assets, are not subject to weekend closures or market holidays.

The block of size 30 balances the bias-variance trade-off inherent to the BMM: smaller blocks yield more observations but weaker extremes, while larger blocks produce purer extremes at the cost of fewer data points [16, 26].

Table 5.1 summarizes the number of observations and resulting blocks for each asset. Bitcoin presents the largest dataset with 188 blocks, followed by Ethereum (127 blocks) and Tether (107 blocks). The dataset for USD Coin (88 blocks) remains at the lower boundary of what can be considered adequate for reliable GEV estimation. This is the fundamental limitation of the BMM: regardless of the total number of observations available, the method retains only one maximum and one minimum per block, so that the entire estimation is conducted on a drastically reduced set of observations.

Table 5.1: Number of daily observations and resulting monthly blocks (block size = 30 days) for each cryptocurrency.

Asset	**Observations**	**Blocks**
Bitcoin (BTC)	5,661	188
Ethereum (ETH)	3,853	127
Tether (USDT)	3,221	107
USD Coin (USDC)	2,574	88

5.1.1 GEV Parameter Estimates

Using the `fgev` function [1], the GEV distribution is fitted to the block maxima of both the loss and gain series for each asset. Parameter estimates are obtained via maximum likelihood, along with their standard errors and 95% confidence intervals computed through the normal approximation [67]. Profile likelihood confidence intervals are additionally derived for the shape parameter ξ to ensure more reliable inference [16]. The results are presented for each asset individually.

Bitcoin (BTC) Table 5.2 reports the GEV parameter estimates for the BTC loss and gain series, together with standard errors.

Table 5.2: GEV parameter estimates and corresponding standard errors for the Bitcoin Block Maxima.

	Loss Series		**Gain Series**	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
μ	1.9339	0.0632	1.8756	0.0615
σ	0.7791	0.0457	0.7458	0.0447
ξ	0.0061	0.0480	-0.0239	0.0558

For the loss series, the GEV maximum likelihood estimates are $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (1.9339, 0.7791, 0.0061)$. The mildly positive shape parameter is suggestive of heavy-tailed behavior consistent with the Fréchet family, though it lies close enough to zero to also be consistent with a Gumbel distribution. The 95% normal approximation confidence interval $[-0.088, 0.100]$ includes zero, indicating that this distinction is not statistically significant, a conclusion further supported by the profile likelihood interval $[-0.0762, 0.1106]$ (Figure 5.1). Hence, the data do not allow one to distinguish between a Fréchet and a Gumbel-type tail for extreme losses.

For the gain series, the estimates are $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (1.8756, 0.7458, -0.0239)$ with a mildly negative shape parameter, placing the distribution within the Weibull family and suggesting a bounded upper tail. However, the corresponding confidence intervals $[-0.1334, 0.0855]$ (normal approximation) and $[-0.1232, 0.0927]$ (profile likelihood, Figure 5.1) again, as in the losses case, straddle zero, providing no statistically significant evidence to support this conclusion.

Overall, despite the relatively large sample size, neither tail yields a shape parameter significantly different from zero. This suggests that, after GARCH filtering, the BTC residuals are consistent with a Gumbel-type distribution, implying exponentially decaying tails rather than the polynomial decay characteristic of the Fréchet family.

Ethereum (ETH) Table 5.3 reports the GEV parameter estimates for the ETH loss and gain series, together with standard errors.

For the loss series, the estimates are $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (1.9689, 0.7206, -0.0818)$ yield a negative shape parameter close to zero, pointing toward the Weibull family. Both the 95% normal approximation confidence interval $[-0.1863, 0.0227]$ and the profile likelihood interval $[-0.1684, 0.0404]$ contain zero (Figure 5.2), providing no statistically significant evidence of heavy-tailed behavior in the ETH loss series.

[1] https://www.rdocumentation.org/packages/evd/versions/2.3-7.1/topics/fgev

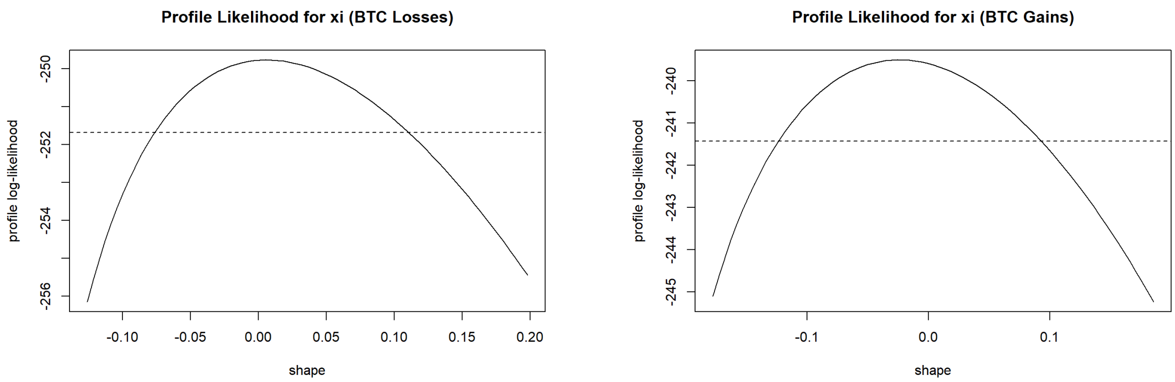

(a) Loss series

(b) Gain series

Figure 5.1: Profile log-likelihood for the shape parameter ξ of the GEV fit to the Bitcoin Block Maxima. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

Table 5.3: GEV parameter estimates and corresponding standard errors for the Ethereum Block Maxima.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
μ	1.9689	0.0706	1.7223	0.0656
σ	0.72061	0.0498	0.6596	0.0475
ξ	-0.08177	0.0533	0.0065	0.0620

For the gain series, the estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (1.7223, 0.6596, 0.0065)$ in the ξ parameter correspond to a Fréchet type of tail. The 95% normal approximation confidence interval $[-0.1151, 0.1282]$ and the profile likelihood interval $[-0.0997, 0.1445]$ both straddle zero (Figure 5.2), indicating that the data do not provide sufficient evidence to actually distinguish between a Fréchet and a Gumbel-type tail.

Ethereum thus mirrors Bitcoin in exhibiting no statistically conclusive evidence of heavy-tailed behavior in either tail.

Tether (USDT) Table 5.4 reports the GEV parameter estimates for the USDT loss and gain series, together with standard errors. For the loss series, the estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (2.0865,$

Table 5.4: GEV parameter estimates and corresponding standard errors for the Tether Block Maxima.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
μ	2.0865	0.1069	2.0652	0.1097
σ	0.9822	0.1002	1.0237	0.0932
ξ	0.4629	0.0879	0.3244	0.0723

$0.9822, 0.4629)$ yield a clearly positive shape parameter, with the 95% normal approximation confidence interval $[0.2907, 0.6352]$ lying entirely above zero, providing strong evidence of Fréchet behavior. The profile likelihood interval $[0.3065, 0.6525]$ (Figure 5.4) is fully consistent with this conclusion. For the gain series, the estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (2.0652, 1.0237, 0.3244)$ tell the same story: both the 95% normal approximation confidence interval $[0.1826, 0.4661]$

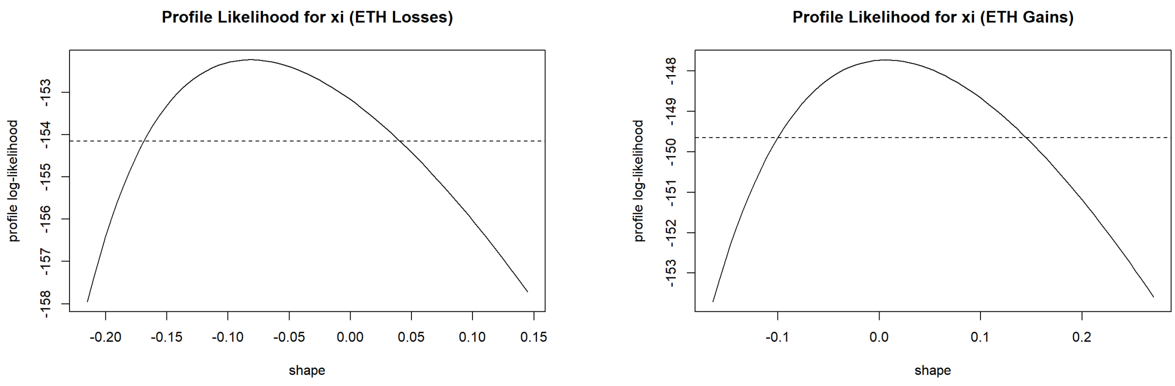

(a) Loss series

(b) Gain series

Figure 5.2: Profile log-likelihood for the shape parameter ξ of the GEV fit to the Ethereum Block Maxima. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

and the profile likelihood interval $[0.1976, 0.4819]$ (Figure 5.4) are strictly positive, confirming statistically significant heavy-tailed Fréchet behavior in both tails.

USD Coin (USDC) Table 5.5 reports the GEV parameter estimates for the USDC loss and gain series, together with standard errors and 95% normal approximation confidence intervals.

Table 5.5: GEV parameter estimates and corresponding standard errors for the USD Coin Block Maxima.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
μ	1.7909	0.0987	1.9829	0.1075
σ	0.8347	0.0769	0.9279	0.0897
ξ	0.1857	0.0738	0.3115	0.0654

For the loss series, the estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (1.7909, 0.8347, 0.1857)$ yield a positive shape parameter, placing the distribution within the Fréchet family and indicating heavy-tailed behavior. This is confirmed by the 95% normal approximation confidence interval $[0.0410, 0.3303]$, which lies entirely above zero, and further supported by the profile likelihood interval $[0.0577, 0.3467]$ (Figure 5.3).

For the gain series, the estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi}) = (1.9829, 0.9279, 0.3115)$ point to an even more pronounced Fréchet tail. The 95% normal approximation confidence interval $[0.1834, 0.4397]$ and the profile likelihood interval $[0.1993, 0.4562]$ (Figure 5.3) are both strictly positive, providing strong and robust evidence that extreme gains in USDC follow a heavy-tailed distribution.

Compared to Bitcoin and Ethereum, both stablecoins exhibit substantially heavier tails, as reflected in the notably larger shape parameter estimates.

5.1.2 Return Level Estimates

When analyzing tail behaviour, a key quantity of interest is the return level, which provides a direct measure of potential portfolio risk in a financial context. As introduced in Section 2.2.1,

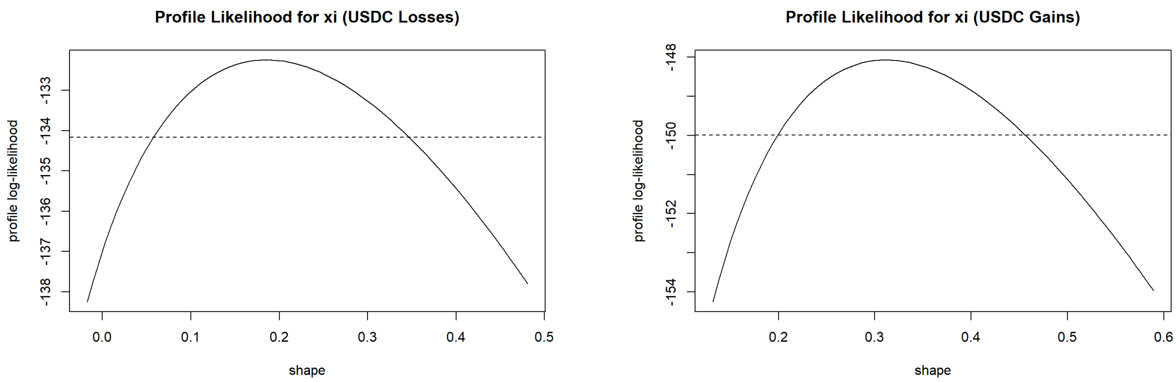

(a) Loss series (b) Gain series

Figure 5.3: Profile log-likelihood for the shape parameter ξ of the GEV fit to the USDC Coin Block Maxima. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

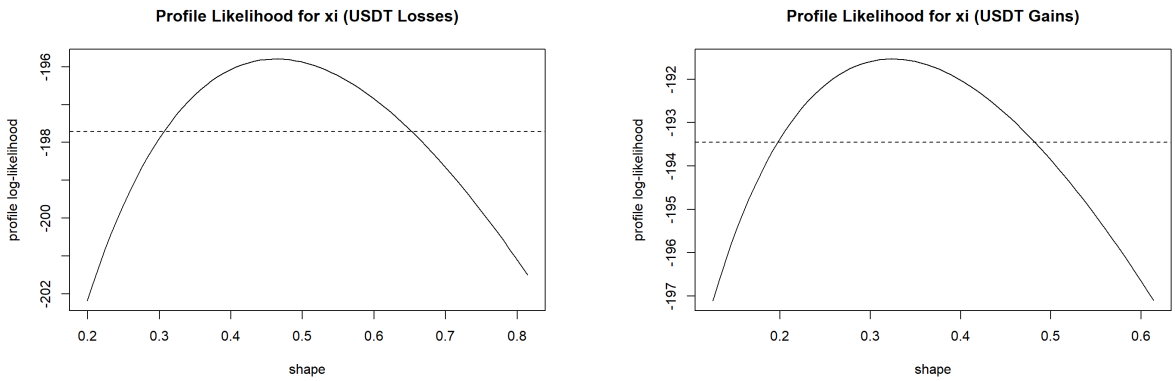

(a) Loss series (b) Gain series

Figure 5.4: Profile log-likelihood for the shape parameter ξ of the GEV fit to the Tether Block Maxima. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

the p-block return level z_p is the threshold expected to be exceeded exactly once in p consecutive Block Maxima. For instance, the 60-block return level is the value that exactly one of 60 monthly maxima is expected to surpass.

The choice of $p = 60$, corresponding to a 5-year return period, reflects a practical consideration. As Coles [16] notes, return level estimates lose credibility as the return period grows relative to the length of the data, since the model is asked to extrapolate well into the tail where little or no empirical information exists. A 60-month horizon strikes a balance between financial relevance and statistical prudence.

Since the GEV is fitted to the standardized GARCH residuals, return levels are initially expressed on the *residual scale*, and cannot be interpreted directly as economically meaningful quantities. To recover such quantities, we invert the standardization of Equation 4.2 for the purely GARCH(1,1) models

$$z_p = \hat{c} + \hat{\sigma}_t \, z_{\text{res}},$$

where $\hat{c}$ is the fixed mean defined by the GARCH model, and $\hat{\sigma}_t$ is the last-day conditional volatility, from the GARCH model. Using the most recent estimates ensures that return levels reflect prevailing market conditions. Because this mapping is linear, standard errors and confidence intervals transform accordingly

$$\mathrm{SE}_{\mathrm{ret}} = \hat{\sigma}_t\, \mathrm{SE}_{\mathrm{res}}, \qquad \mathrm{CI}_{\mathrm{ret}} = \hat{c} + \hat{\sigma}_t\, \mathrm{CI}_{\mathrm{res}},$$

so both the delta method and profile likelihood intervals carry over without distortion.

For the ARMA(1,1)-GARCH(1,1), used to filter the stablecoin returns, the return levels are obtained inverting Equation 4.3

$$z_p = \hat{E}_{t-1}[r_t] + \hat{\sigma}_t\, z_{\mathrm{res}},$$

The only difference is that here we must estimate both the last-day conditional mean, $\hat{E}_{t-1}[r_t]$, the conditional standard deviation, $\hat{\sigma}_t$.

Standard errors and confidence intervals, as well as before, transform accordingly

$$\mathrm{SE}_{\mathrm{ret}} = \hat{\sigma}_t\, \mathrm{SE}_{\mathrm{res}}, \qquad \mathrm{CI}_{\mathrm{ret}} = \hat{E}_{t-1}[r_t] + \hat{\sigma}_t\, \mathrm{CI}_{\mathrm{res}},$$

The resulting return-scale estimates are reported in Table 5.6, alongside standard errors and 95% confidence intervals from both methods.

Table 5.6: 60-block (five-year) return levels for losses and gains, rescaled to the log-return scale using the last-day GARCH conditional mean and volatility. Point estimates and standard errors are reported, along with 95% confidence intervals obtained via the delta method and via profile likelihood.

Asset	**Series**	**RL**	**SE**	**Delta Method CI**		**Profile Likelihood CI**	
				Lower	**Upper**	**Lower**	**Upper**
BTC	Loss	1.1404	0.0726	0.9980	1.2828	1.0337	1.3275
BTC	Gain	1.0565	0.0681	0.9231	1.1900	0.9576	1.2299
ETH	Loss	0.8973	0.0511	0.7972	0.9975	0.8278	1.0372
ETH	Gain	0.8936	0.0701	0.7563	1.0309	0.7956	1.0904
USDT	Loss	0.8835	0.2052	0.4813	1.2857	0.6078	0.9432
USDT	Gain	0.6785	0.1195	0.4443	0.9128	0.5104	0.7849
USDC	Loss	3.8419	0.5072	2.8478	4.8361	3.1537	5.3335
USDC	Gain	5.2254	0.8301	3.5986	6.8523	4.0680	5.9824

Before discussing the return level estimates, one clarification is necessary. Since the ARMA–GARCH model for the stablecoins was fitted on returns scaled by a factor of 100, the return level estimates reported in Table 5.6 must be divided by 100 to recover the original log-return scale.

All reported return levels can be mapped back to percentage price declines or increases by reversing the standard log-return transformation introduced in Equation (4.1), namely via $e^{z_p} - 1$. It is important to note that these percentage changes are scale-free: they represent the proportional move relative to whatever the asset's price happens to be on that day, and are therefore directly interpretable as the fraction of the current position that would be lost or gained in a tail event, regardless of the prevailing price level.

Several compelling patterns emerge from Table 5.6, the most striking of which is the sharp contrast between the highly volatile crypto-assets (Bitcoin and Ethereum) and the fiat-pegged stablecoins (USDC and USDT).

For Bitcoin and Ethereum, the estimated return levels are large and economically significant. Bitcoin's five-year conditional loss return level stands at 1.1404, which translates into a severe single-day price decline of $e^{-1.1404} - 1 \approx 68\%$. This represents the conditional worst expected single-day loss to be exceeded exactly once over a five-year horizon. The corresponding gain return level of 1.0565 implies an exceptional single-day price surge of $e^{1.0565} - 1 \approx 187\%$, reflecting the historically violent upside momentum typical of crypto bull markets (such as the 2020–2021 expansion) [11].

Notably, this framework reveals a marked asymmetry in the tails: while the downside crashes are massive, single-day positive rallies at the five-year horizon are even more explosive in proportional terms. This asymmetry is a direct consequence of the positive skewness inherent to speculative cryptocurrency return distributions, a crucial stylized fact that symmetric, Gaussian-based risk models fail to capture. Ethereum displays a qualitatively identical behavior: its five-year loss return level of 0.8973 translates into a price drop of $e^{-0.8973} - 1 \approx 59.2\%$, while its gain return level of 0.8936 corresponds to a price increase of $e^{0.8936} - 1 \approx 144.4\%$, again demonstrating upside tail dominance. For both assets, the standard errors are small and the 95% confidence intervals (both delta method and Profile Likelihood) are narrow, indicating stable tail estimations and reliable asymptotic extrapolations.

The contrast with the stablecoin market is stark. As noted, because the ARMA–GARCH models for the stablecoins were fitted on percentage returns, their return levels must be divided by 100 before applying the exponential transformation. For USDC, the five-year loss return level is adjusted to 0.0384, which implies an expected price decline of $e^{-0.0384} - 1 \approx 3.8\%$. Its gain return level is adjusted to 0.0523, corresponding to a price appreciation of $e^{0.0523} - 1 \approx 5.36\%$. While these figures are orders of magnitude smaller than those of BTC and ETH confirming the stability enforced by the peg mechanism, they are far from negligible. A single-day deviation of approximately 3.8% to 5.36% from the dollar peg for a supposedly stable asset represents a meaningful tail risk. This is consistent with historical episodes such as the brief de-pegging triggered by the Silicon Valley Bank collapse in March 2023 [2].

Finally, USDT presents the most extreme case of tightly bound volatility within the sample. Its adjusted five-year return levels are 0.0088 for losses and 0.0068 for gains, yielding a maximum expected price decline of $e^{-0.0088} - 1 \approx 0.88\%$ and a maximum price increase of $e^{0.0068} - 1 \approx 0.68\%$.

Overall, the conditional GEV framework sharply divide the two asset classes based on their extreme risk profiles. Bitcoin and Ethereum display large, asymmetric five-year horizons, with worst-case single-day losses hovering around 59.23% to 68.03% and worst-case gains reaching 144.39% to 187.63%, fully consistent with the heavy-tailed and volatile nature of unbacked cryptocurrencies. At the same time, the model captures the distinct regimes within the stablecoin space: USDC retains a small but historically justified tail risk of approximately 4% to 5% under stress, whereas USDT, conditional on its recent period of compressed, near-zero volatility, exhibits conditional extreme return levels that are economically minor across the considered horizon.

5.1.3 Model Evaluation

The adequacy of the GEV model fitted via the Block Maxima method is assessed using four standard diagnostic plots [16]. The probability plot compares empirical and theoretical cumulative probabilities, with observations lying close to the 45° line indicating a good fit. The quantile plot is particularly informative in the EVT context, as it emphasizes tail behavior; systematic upward deviations in the upper tail indicate that the model underestimates extreme observations. The return level plot displays the fitted return level curve together with confidence bands and empirical estimates, where observations falling outside the bands, especially at longer return periods, signal model inadequacy in the tail. Recall from the theoretical Chapter 2.2.1 that the shape of the curve also reflects the tail behavior of the underlying distribution: it is linear when $\xi = 0$ (Gumbel case), convex with a finite upper bound $\mu - \sigma/\xi$ when $\xi < 0$, and concave with no finite upper bound when $\xi > 0$, indicating a heavy-tailed distribution. Finally, the density plot provides a visual comparison between the fitted GEV density and the empirical distribution of Block Maxima. These diagnostics are interpreted jointly to assess model adequacy. For brevity, only the most informative plots are presented here, while the remaining ones are reported in Appendix A.

Bitcoin (BTC) The GEV model provides a good overall fit for Bitcoin. The probability plot shows observations closely aligned with the diagonal, and the quantile plot indicates a strong agreement between empirical and theoretical quantiles across most of the distribution, with only minor deviations in the upper tail. The return level plot is broadly consistent, with the majority of Block Maxima falling within the confidence bands. The density plot further supports this assessment, with the fitted distribution closely matching the empirical histogram for both losses and gains. While the fit for gains is marginally weaker than for losses, the Block Maxima approach offers an adequate representation of extreme behavior for this asset.

Ethereum (ETH) A similar conclusion holds for Ethereum. Both the probability and quantile plots display no clear evidence of systematic misspecification, and the model captures tail behavior reasonably well in both series. Some deviation is visible in the upper tail of the quantile plot, and a small number of observations fall outside the return level confidence bands, particularly for gains. However, these discrepancies are limited in magnitude and do not materially affect the overall adequacy of the model.

USD Coin (USDC) Figure 5.5 displays the GEV diagnostic plots for the ARMA–GARCH standardized residuals of USDC gains. While the probability plot suggests a reasonable fit in the central region of the distribution, the quantile and return level plots reveal a severe upward deviation in the extreme upper tail. Specifically, a single macro-outlier stands out dramatically above the fitted GEV curve, lying far outside the 95% confidence band at longer return periods. This phenomenon illustrates a well-known structural limitation of the GARCH framework: its inability to anticipate sudden, instantaneous price jumps (such as unexpected de-pegging events). Because the GARCH model updates its volatility forecasts based on lagged historical information, the conditional variance (σ_t) estimated for the day of the shock was too low, failing to properly standardize and "shrink" the extreme return [51]. A less severe but qualitatively similar pattern emerges for USDC losses. The density and probability plots both show a cleaner fit compared to the gains, with the distribution tracking the theoretical curve more faithfully across most of its range. Still, the quantile and return level plots reveal that one outlier manages to slip outside the 95% confidence band, the same structural issue as before, just less dramatic. The point sits just past the edge of the confidence interval, suggesting that the extreme loss event, while still outside the model's reach, was comparatively less disruptive to the overall fit.

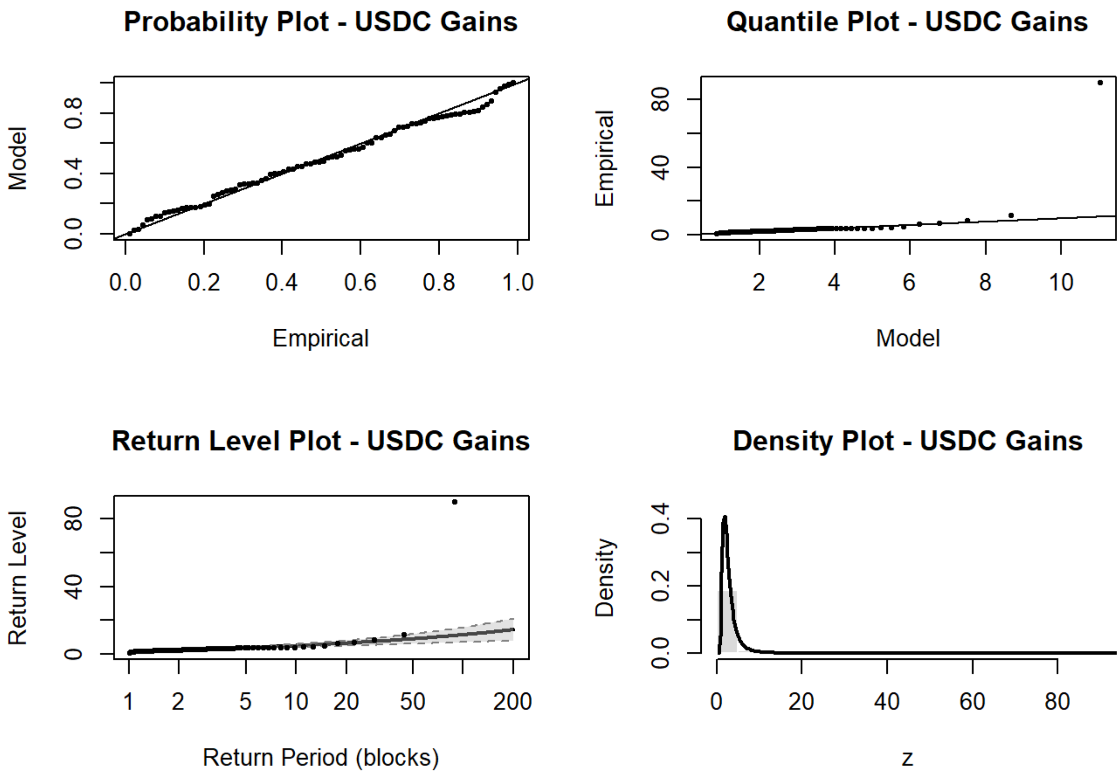

Figure 5.5: GEV diagnostic plots for USDC gains. The quantile and return level plots reveal a severe upward deviation in the extreme upper tail, with a single macro-outlier lying far outside the 95% confidence band at longer return periods.

Tether (USDT) The diagnostic plots for the USDT gains residuals, shown in Figure 5.6, indicate once again a lack of fit. The probability plot displays some deviations from the diagonal, but the most striking feature is the severe departure in the quantile and return level plots, where an isolated outlier breaks completely away from the fitted GEV distribution, breaching the 95% confidence bands at longer horizons, mirroring the USDC gains case. A similar conclusion holds for USDT losses, where deviations are again visible in the probability plot and the density plot also fails to achieve a satisfactory fit. The quantile and return level plots likewise display suboptimal results. This time, however, no particular outlier dramatically exits the 95% confidence band, yet the overall fit remains far from optimal.

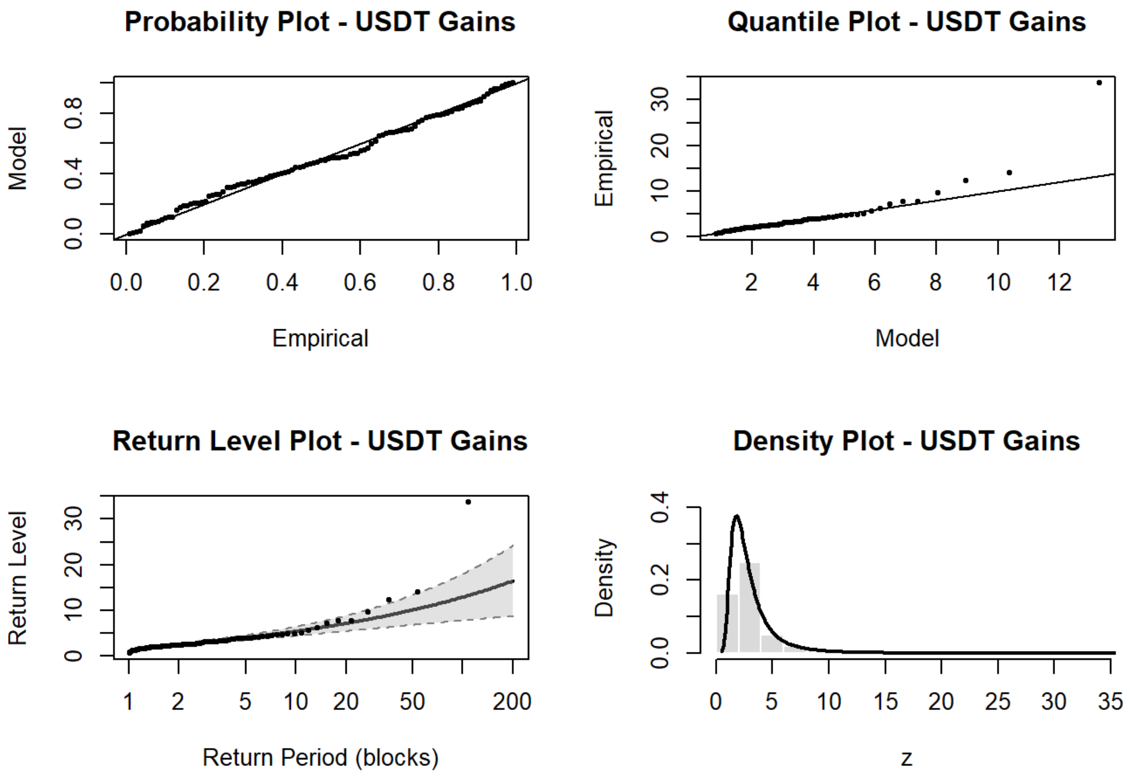

Figure 5.6: GEV diagnostic plots for USDT gains. The quantile and return level plots reveal a severe departure from the fitted GEV distribution, with an isolated outlier breaching the 95% confidence band at longer return periods, mirroring the pattern observed for USDC gains.

Overall, the diagnostic analysis reveals a clear separation between the two asset classes. For USDC and USDT, the GEV model fitted via the Block Maxima approach fails to fully capture the behavior of the stablecoins, as evidenced by consistent deviations in the quantile and return level plots. It is worth noting that stablecoins are, by their very nature, among the most difficult assets to model: for the vast majority of the time they exhibit extremely low variation, only to experience sudden spikes that are almost immediately corrected through the mechanisms that support the peg, producing a return distribution that is highly atypical and particularly challenging to capture with standard Extreme Value tools. Their relatively shorter time series, combined with the limited number of observations retained by the BMM method, further compounds this challenge, reducing the precision of tail estimation.

In contrast, the model provides a substantially better fit for BTC and ETH, consistent with their near-Gumbel shape parameter estimates, which imply lighter and more stable tails. These findings highlight a fundamental limitation of the Block Maxima approach in this setting: by construction, it discards the majority of extreme observations, leading to reduced efficiency and potential underestimation of tail risk. These limitations motivate the adoption of the Peaks Over Threshold method in the following section, which exploits all exceedances above a high threshold and is therefore expected to yield a more efficient and reliable characterization of extreme risk across all assets.

5.2 Peaks Over Threshold Approach

5.2.1 Threshold Selection

The first step in the Peaks-Over-Threshold approach is to select an appropriate threshold u for each series. This choice reflects the bias-variance trade-off inherent in threshold selection. In empirical implementation, low thresholds introduce contamination from non-tail observations, whereas high thresholds reduce the number of exceedances and increase estimation uncertainty.

Following the two-step procedure described in Section 2.3, the threshold is determined by jointly examining the mean residual life plot and the parameter stability analysis for each series. Specifically, the selected threshold is the lowest value of u at which the shape parameter ξ and the modified scale parameter $\sigma^* = \tilde{\sigma}_u - \xi u$ no longer exhibit systematic trends and instead fluctuate within their respective confidence intervals, in line with the criterion of [16]. The mean residual life plots for all series are shown in Figure 5.7 for gains and Figure 5.8 for losses, with the red dashed vertical line indicating the chosen threshold based on the stability analysis.

The selected thresholds are summarized in Table 5.7, together with the corresponding number of exceedances available for estimation. Two features of this table are worth noting. First, the stablecoins USDC and USDT require substantially higher thresholds than Bitcoin and Ethereum, reflecting the heavy Fréchet tail behavior identified in the Block Maxima analysis: for these assets, the GPD approximation only becomes valid deeper into the tail, and full parameter stabilization is not always achieved, consistent with the presence of an extremely heavy tail. For BTC and ETH, by contrast, stability is attained at low thresholds, consistent with the near-Gumbel shape parameter estimates that imply a well-behaved and rapidly stabilizing tail. Second, and importantly, the number of exceedances available for GPD estimation is considerably larger than the number of Block Maxima retained in the GEV analysis. This is a direct consequence of the greater efficiency of the POT approach, which conditions on all observations exceeding the threshold rather than discarding all but the single maximum within each block. The increased sample size is expected to translate into more precise parameter estimates and more reliable return level extrapolations, addressing one of the key limitations of the Block Maxima method identified in the previous section.

Table 5.7: Selected thresholds and number of exceedances for each series.

Asset	**Side**	**Threshold** u	**Exceedances** n_u
BTC	Losses	0.78	877
	Gains	0.86	714
ETH	Losses	0.78	591
	Gains	0.75	575
USDT	Losses	1.40	255
	Gains	1.35	266
USDC	Losses	1.45	219
	Gains	1.26	173

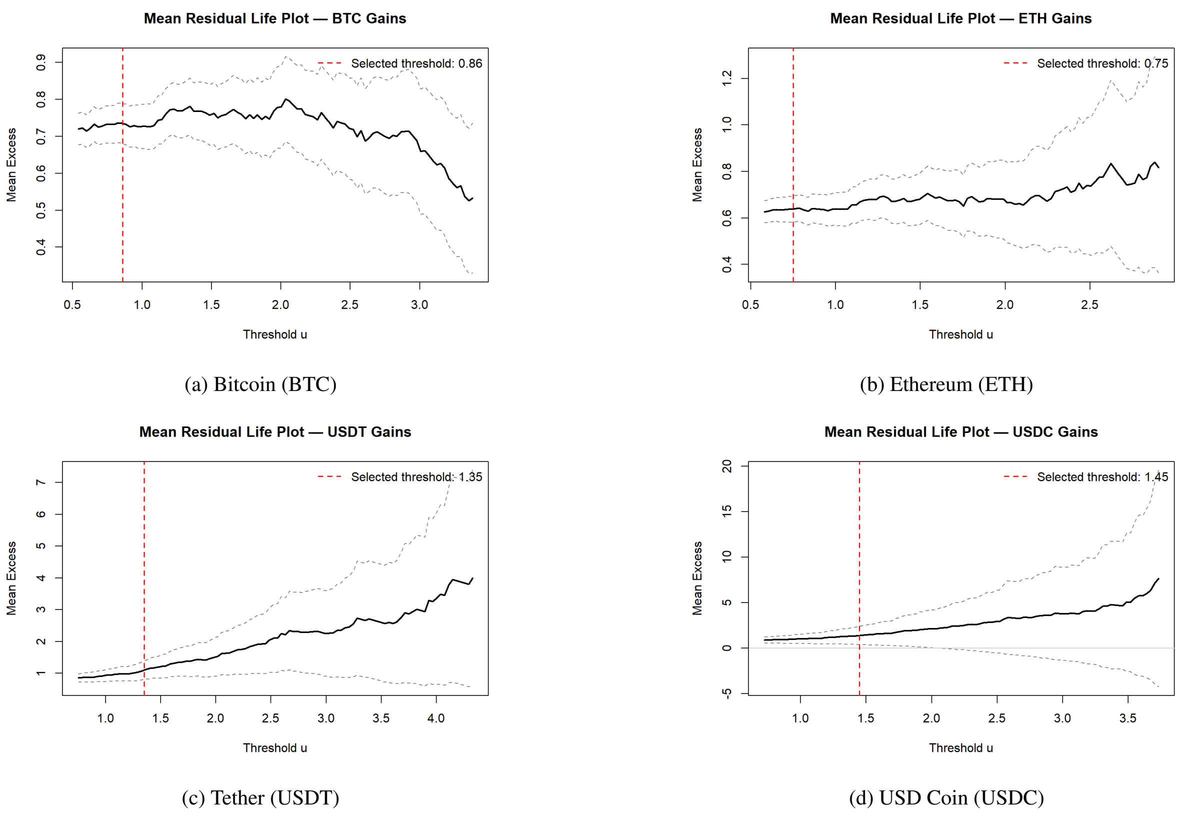

(a) Bitcoin (BTC)

(b) Ethereum (ETH)

(c) Tether (USDT)

(d) USD Coin (USDC)

Figure 5.7: Mean residual life plots for gains across all four assets. The red dashed vertical line indicates the threshold selected on the basis of the parameter stability analysis.

5.2.2 GPD Parameter Estimates

The Generalized Pareto Distribution is fitted using the `fpot` function [2], separately to the exceedances above the selected threshold for the loss series and the gain series of each cryptocurrency, yielding estimates of the scale parameter $\tilde{\sigma}$ and the shape parameter ξ. Parameter estimates are obtained by maximum likelihood estimation, with standard errors and approximate 95% confidence intervals derived via the normal approximation. To obtain more accurate inference on ξ specifically, given that the asymptotic normality assumption underlying the normal approximation may be unreliable for tail index estimation in moderate samples [16], profile likelihood confidence intervals are additionally computed for the shape parameter of each series. The results for each asset are discussed in turn.

Bitcoin (BTC) Table 5.8 reports the GPD parameter estimates for the BTC loss and gain series, together with standard errors and 95% normal approximation confidence intervals.

For the loss series, maximum likelihood estimation yields estimates $(\hat{\tilde{\sigma}}, \hat{\xi}) = (0.6825, 0.0497)$. The shape parameter is greater than 0, corresponding to a Fréchet distribution. The 95% normal approximation confidence interval for ξ, equal to $[-0.0199, 0.1192]$, and the profile likelihood interval $[-0.01397, 0.12489]$, fully consistent with the result of Figure 5.9, do not discard a Gumbel family since they contain 0. There is no statistically meaningful evidence of heavy-tailed behavior in the BTC loss series.

For the gain series, the situation is very similar: the estimates are $(\hat{\tilde{\sigma}}, \hat{\xi}) = (0.7330, 0.0043)$

[2] https://www.rdocumentation.org/packages/evd/versions/2.3-7.1/topics/fpot

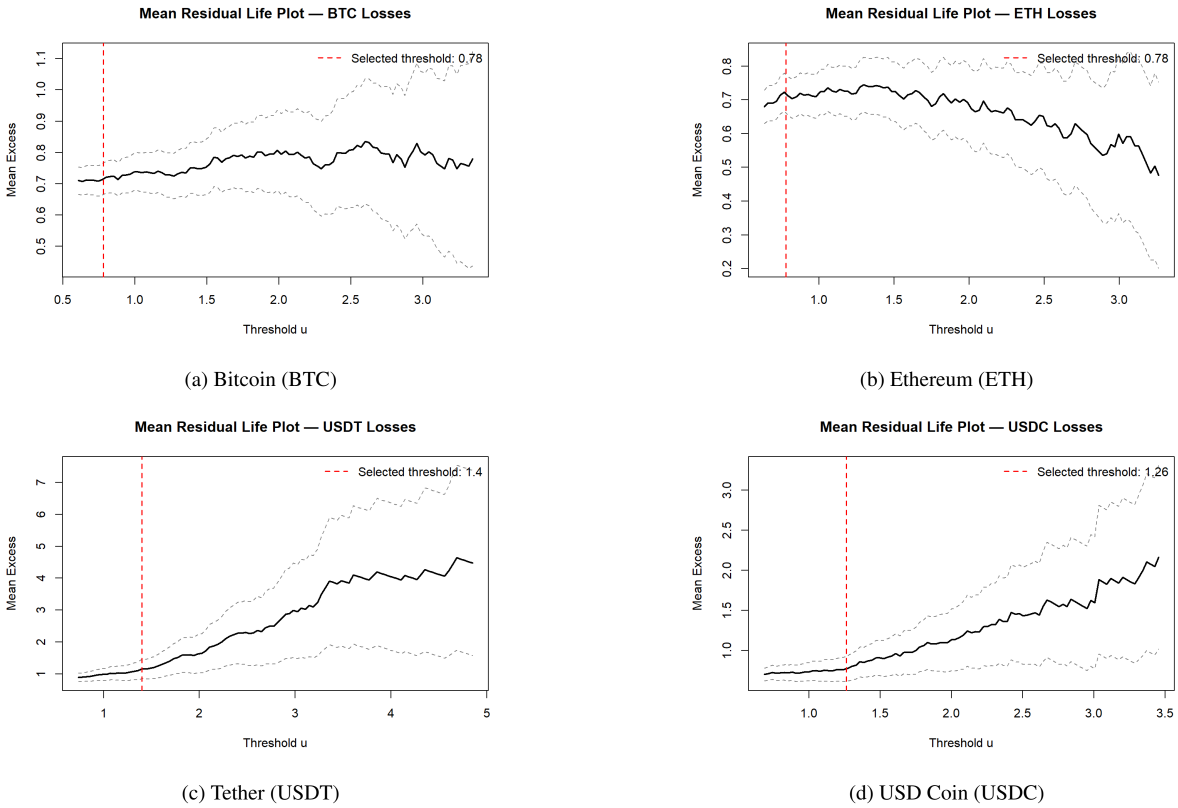

(a) Bitcoin (BTC)

(b) Ethereum (ETH)

(c) Tether (USDT)

(d) USD Coin (USDC)

Figure 5.8: Mean residual life plots for losses across all four assets. The red dashed vertical line indicates the threshold selected on the basis of the parameter stability analysis.

Table 5.8: GPD parameter estimates for Bitcoin (BTC) exceedances above $u = 0.78$ (losses) and $u = 0.86$ (gains), with standard errors.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
$\tilde{\sigma}$	0.6825	0.0334	0.7330	0.0407
ξ	0.0497	0.0355	0.0043	0.0409

and again the shape parameter is slightly positive, but the 95% normal approximation confidence interval $[-0.0760, 0.0845]$ straddles zero, and the profile likelihood interval $[-0.0694, 0.0908]$ corroborates this result (Figure 5.9). As with the loss series, the evidence is insufficient to conclusively establish heavy-tailed behavior in BTC extreme gains.

Bitcoin therefore exhibits no statistically conclusive evidence of heavy-tailed behavior in either tail under the POT framework, with shape parameter estimates close to zero and confidence intervals straddling zero under both inference methods. This is fully consistent with the near-Gumbel findings from the Block Maxima analysis, and suggests that the tail of the BTC standardized residual distribution decays exponentially rather than polynomially.

Ethereum (ETH) Table 5.9 reports the GPD parameter estimates for the ETH loss and gain series, together with standard errors and 95% normal approximation confidence intervals.

For the loss series, maximum likelihood estimation yields estimates $(\hat{\tilde{\sigma}}, \hat{\xi}) = (0.7265, -0.0177)$. The shape parameter estimate is negative, suggesting a Weibull-type distribution, but the 95% normal approximation confidence interval of $[-0.1025, 0.0671]$ straddles zero, so the Gumbel

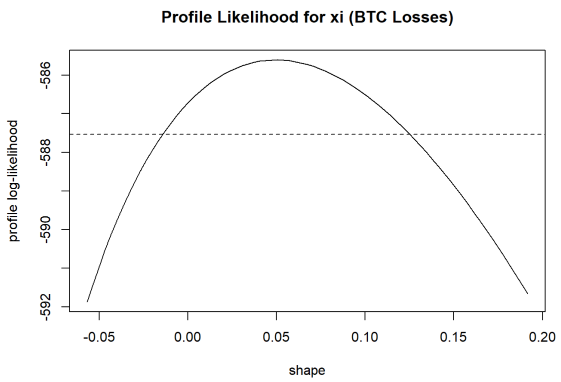

(a) Loss series

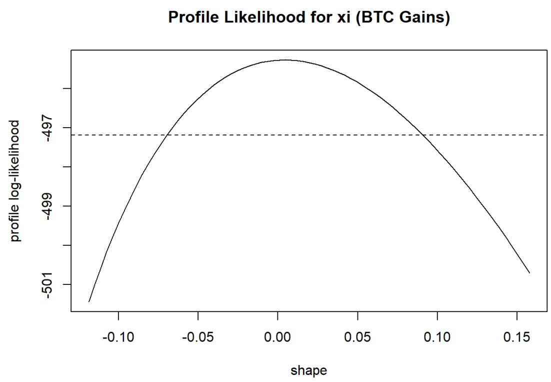

(b) Gain series

Figure 5.9: Profile log-likelihood for the shape parameter ξ of the POT fit to Bitcoin exceedances. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

Table 5.9: GPD parameter estimates for Ethereum (ETH) exceedances above $u = 0.78$ (losses) and $u = 0.75$ (gains), with standard errors.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
$\tilde{\sigma}$	0.7265	0.0434	0.5991	0.0356
ξ	-0.0177	0.0433	0.0578	0.0425

family cannot be rejected. The profile likelihood interval of $[-0.0936, 0.0756]$ is consistent with this finding (Figure 5.10).

For the gain series, the estimates are $(\hat{\tilde{\sigma}}, \hat{\xi}) = (0.5991, 0.0578)$. The shape parameter point estimate is small but positive, suggesting a Fréchet-type distribution. The 95% normal approximation confidence interval of $[-0.0256, 0.1411]$ straddles zero, and the profile likelihood interval of $[-0.0169, 0.1499]$ corroborates this result (Figure 5.10). The data therefore do not provide sufficient evidence to distinguish between the Fréchet and Gumbel families for ETH extreme gains.

Ethereum thus mirrors Bitcoin in exhibiting no statistically conclusive evidence of heavy-tailed behavior in either tail, with shape parameter estimates close to zero and confidence intervals straddling zero under both inference methods. These findings are consistent with the Block Maxima results and reinforce the conclusion that ETH standardized residuals are well described by a near-Gumbel tail.

Tether (USDT) Table 5.10 reports the GPD parameter estimates for the USDT loss and gain series, together with standard errors and 95% normal approximation confidence intervals.

Table 5.10: GPD parameter estimates for Tether (USDT) exceedances above $u = 1.40$ (losses) and $u = 1.35$ (gains), with standard errors.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
$\tilde{\sigma}$	0.6056	0.0618	0.6252	0.0639
ξ	0.4669	0.0857	0.4186	0.0858

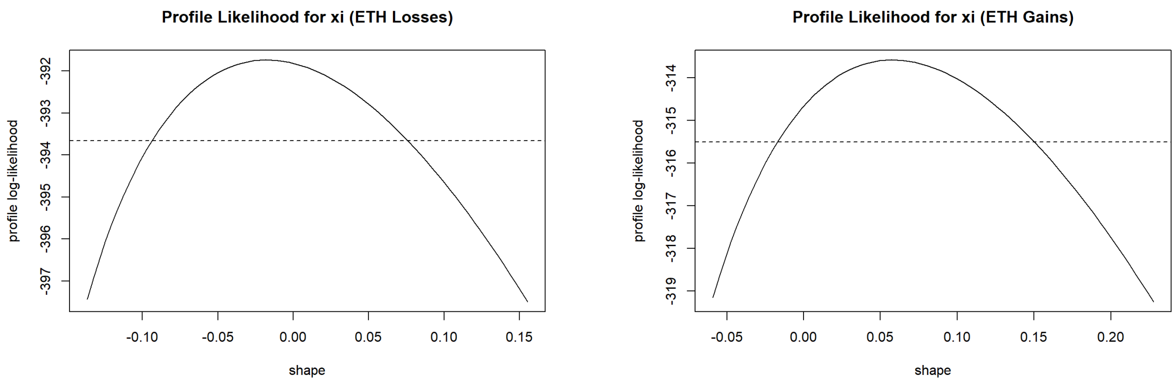

(a) Loss series

(b) Gain series

Figure 5.10: Profile log-likelihood for the shape parameter ξ of the POT fit to Ethereum exceedances. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

For the loss series, maximum likelihood estimation yields estimates $(\hat{\sigma}, \hat{\xi}) = (0.6056, 0.4669)$. The shape parameter estimate is the largest among all series considered, and the 95% normal approximation confidence interval of $[0.2990, 0.6348]$ lies entirely above zero, providing strong evidence of heavy-tailed Fréchet-type behavior. The profile likelihood interval of $[0.3171, 0.6553]$ confirms this result (Figure 5.11), and its asymmetric shape reflects the pronounced positive skewness of the likelihood surface for large values of ξ.

For the gain series, the estimates are $(\hat{\sigma}, \hat{\xi}) = (0.6252, 0.4186)$. The shape parameter is again positive, with a 95% normal approximation confidence interval of $[0.2504, 0.5867]$ lying entirely above zero. The profile likelihood interval of $[0.2701, 0.6084]$ corroborates this result (Figure 5.11). USDT therefore exhibits statistically significant heavy-tailed behavior in both tails, consistent with the GEV findings and further supported by the increased precision afforded by the POT approach.

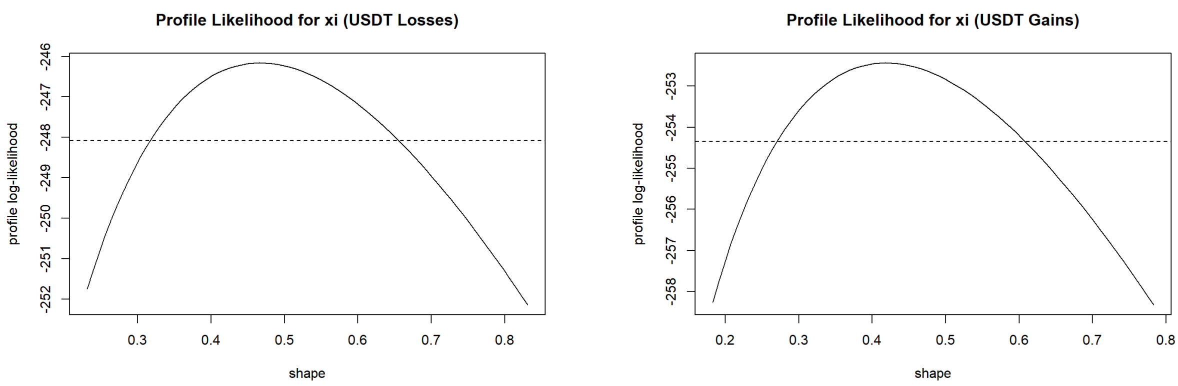

(a) Loss series

(b) Gain series

Figure 5.11: Profile log-likelihood for the shape parameter ξ of the POT fit to Tether exceedances. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

USD Coin (USDC) Table 5.11 reports the GPD parameter estimates for the USDC loss and gain series, together with standard errors and 95% normal approximation confidence intervals.

Table 5.11: GPD parameter estimates for USD Coin (USDC) exceedances above $u = 1.26$ (losses) and $u = 1.45$ (gains), with standard errors.

	Loss Series		Gain Series	
Parameter	**Estimate**	**SE**	**Estimate**	**SE**
$\tilde{\sigma}$	0.5071	0.0568	0.6162	0.0741
ξ	0.3550	0.0926	0.4262	0.0986

For the loss series, maximum likelihood estimation yields estimates $(\hat{\tilde{\sigma}}, \hat{\xi}) = (0.5071, 0.3550)$. The shape parameter estimate is positive and the 95% normal approximation confidence interval of $[0.1735, 0.5364]$ lies entirely above zero, providing statistically significant evidence of heavy-tailed Fréchet-type behavior in the USDC loss series. The profile likelihood interval of $[0.1959, 0.5616]$ corroborates this finding (Figure 5.12).

For the gain series, the estimates are $(\hat{\tilde{\sigma}}, \hat{\xi}) = (0.6162, 0.4262)$. The shape parameter is more decisively positive, with a 95% normal approximation confidence interval of $[0.2330, 0.6194]$ lying entirely above zero. The profile likelihood interval of $[0.2633, 0.6532]$ confirms this result (Figure 5.12), reinforcing the conclusion that extreme gains in USDC exhibit strongly heavy-tailed behavior under the POT framework. These findings are consistent with the corresponding GEV results, suggesting that both methods successfully characterize the tail behavior.

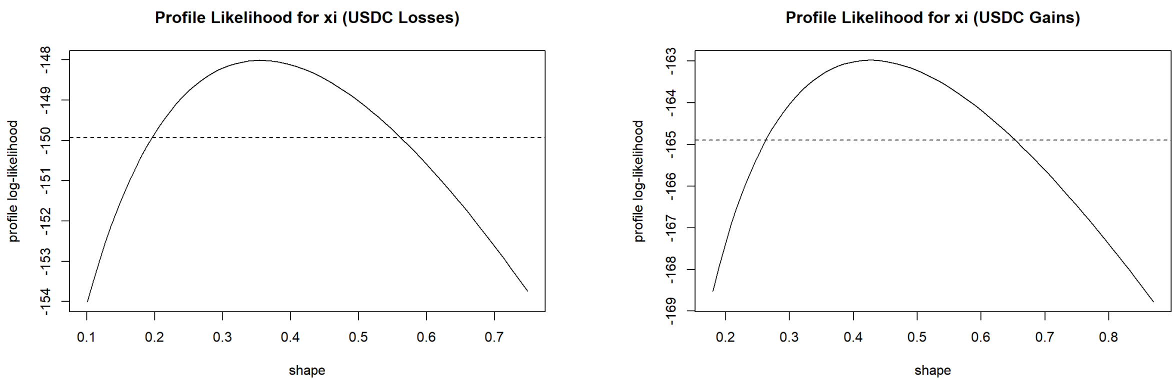

(a) Loss series

(b) Gain series

Figure 5.12: Profile log-likelihood for the shape parameter ξ of the POT fit to USDC Coin exceedances. The dashed horizontal line indicates the 95% critical threshold, from which the profile likelihood confidence interval is derived.

Overall, the GPD parameter estimates are in close agreement with those obtained under the GEV framework, not only in terms of tail classification, but also in the scale estimates and the sharp contrast between volatile assets and stablecoins.

5.2.3 Return Level Estimates

Having established the GPD parameter estimates for all four assets, we now turn to the estimation of return levels under the Peaks-Over-Threshold framework. The m-observation return level x_m represents the level exceeded on average once every m observations, where m is expressed in trading days. Return levels are estimated at the 1260-observation horizon, corresponding to a five-year return period, for both the loss and gain series of each asset. As in the Block Maxima section, the estimated return levels are converted back from the filtered standardized residual scale to the original daily log-return scale. Following the framework of McNeil and Frey [51]. Practically, this is performed scaling the static GPD quantiles using the last day conditional volatility and mean values from the ARMA-GARCH model for the stablecoins, and the last day conditional volatility value for the pure GARCH models (where the conditional mean is fixed). This step ensures that all final figures are directly interpretable as actual daily log-returns.

Table 5.12: Estimated five-year (1260-observation) return levels for losses and gains of each asset, expressed as daily log-returns after rescaling from the GARCH-filtered standardized residual scale using the last-day conditional mean and volatility. Point estimates and standard errors are reported, along with 95% confidence intervals obtained via the delta method and via profile likelihood.

Asset	**Series**	**RL**	**SE**	**Delta Method CI**		**Profile Likelihood CI**	
				Lower	**Upper**	**Lower**	**Upper**
BTC	Loss	1.0824	0.0665	0.9520	1.2127	0.9803	1.2389
BTC	Gain	1.0205	0.0612	0.9005	1.1404	0.9240	1.1680
ETH	Loss	0.8905	0.0583	0.7762	1.0049	0.8066	1.0381
ETH	Gain	0.8871	0.0655	0.7588	1.0154	0.7909	1.0528
USDT	Loss	0.7063	0.1404	0.4311	0.9815	0.5145	1.1334
USDT	Gain	0.6474	0.1228	0.4068	0.8881	0.4806	1.0208
USDC	Loss	4.0224	0.6676	2.7139	5.3309	3.1385	6.1428
USDC	Gain	5.1347	0.9745	3.2247	7.0448	3.8477	8.3693

The results in Table 5.12 reveal a stark contrast between the volatile assets and the stablecoins, fully consistent with the GEV findings of Table 5.6. As for the GEV case all the return levels can be interpreted as price decline and increase, reversing the definition of the return levels, as introduced in 4.1.

For Bitcoin, the five-year loss return level is estimated at 1.0824, which corresponds to a price decline of $e^{1.0824} - 1 \approx 66.1\%$, while the gain return level stands at 1.0205, translating into a price increase of $e^{1.0205} - 1 \approx 177.5\%$. These figures are in line with those obtained under the GEV framework, providing a sign of robustness across both EVT methodologies. Ethereum exhibits slightly lower but highly comparable figures: the five-year loss return level of 0.8905 translates into a price decline of $e^{0.8905} - 1 \approx 58.9\%$, and the gain return level of 0.8871 implies a price increase of $e^{0.8871} - 1 \approx 142.8\%$.

When analyzing the stablecoins, a first crucial step is required. As for the GEV case, the values reported in the table for USDC and USDT must be scaled down by a factor of 100 to account for the return adjustments. For USDC, the true five-year loss return level is 0.0402, which corresponds to a modest price decline of $e^{0.0402} - 1 \approx 3.9\%$. The gain return level is 0.0513, translating into a price increase of $e^{0.0513} - 1 \approx 5.3\%$, that is really close to the $\approx 5.36\%$ decline estimated in the GEV. These figures are roughly two orders of magnitude smaller than those of the volatile assets, confirming that USDC extreme moves remain small in absolute terms. As noted in the GEV discussion, a 4%-5% extreme move reflects the occasional stress episodes that even asset-backed pegs can experience during periods of severe market-wide liquidity crunches.

USDT presents an even more contained case, with the scaled five-year return levels dropping to 0.0071 for losses and 0.0065 for gains. These values translate into a price decline of $e^{-0.0071} - 1 \approx 0.70\%$ and a price increase of $e^{0.0065} - 1 \approx 0.65\%$, respectively. This outcome is fully consistent with the near-zero tail risk found in the GEV analysis. As previously noted, this behavior comes from how the GPD model and the GARCH filter interact. Because USDT returns were extremely stable during the final part of the sample, the GARCH model predicts a very low conditional volatility. When the GPD quantiles are multiplied by this tiny scale parameter, the final return levels collapse toward zero.

Value-at-Risk and Expected Shortfall

Table 5.13 reports the Value-at-Risk and Expected Shortfall for the loss series of each asset at two confidence levels. The 99% confidence level corresponds to a loss expected to be exceeded once every 100 trading days, while the 99.92% level corresponds to once every 1260 trading days, precisely the five-year horizon adopted throughout this work. By construction, the $\mathrm{VaR}_{99.92\%}$ is therefore directly comparable to the five-year return levels reported in Table 5.12.

Table 5.13: Value-at-Risk and Expected Shortfall at the 99% and 99.92% confidence levels for the loss series of each asset, expressed as daily log-returns after rescaling from the GARCH-filtered standardized residual scale using the last-day conditional mean and volatility. The 99.92% level corresponds to the five-year return horizon and is directly comparable to the return levels reported in Table 5.12.

	99%		**99.92%**	
Asset	**VaR**	**ES**	**VaR**	**ES**
BTC	0.6149	0.7971	1.0808	1.2874
ETH	0.5462	0.6827	0.8895	1.0201
USDT	0.2179	0.4069	0.7037	1.3180
USDC	1.8032	2.6408	4.0118	6.0648

For Bitcoin and Ethereum the results are consistent with the return level estimates. The BTC $\mathrm{VaR}_{99\%}$ of 0.6149 corresponds to a price decline of $e^{-0.6149} - 1 \approx 45.9\%$, representing a loss threshold exceeded roughly once every 100 trading days. The corresponding ES of 0.7971, $e^{0.7971} - 1 \approx 54.94\%$, indicates that on days when losses surpass this threshold, the average loss is substantially larger. At the five-year horizon, the $\mathrm{VaR}_{99.92\%}$ of 1.080 aligns closely with the return level of 1.0824 from Table 5.12, as expected by construction, and on days when losses surpass the threshold, the average loss is $e^{-1.2874} - 1 \approx 72.4\%$.

Ethereum follows the same pattern. The $\mathrm{VaR}_{99\%}$ of 0.5462 corresponds to a price decline of $\approx 42\%$, with an ES of 0.6827 indicating that average losses on exceedance days are substantially larger than the threshold, and are $\approx 49.5\%$. The five-year $\mathrm{VaR}_{99.92\%}$ of 0.8895 is consistent with the corresponding return level of 0.8905, corresponding to a price decline of $\approx 59\%$ beyond which the average tail loss progresses to $e^{-1.0201} - 1 \approx -63.94\%$.

The values recorded for USDC, as noted in previous sections, must be divided by 100 to reverse the original operational rescaling and allow for an accurate interpretation on the proper daily log-return scale. At the 99% confidence level, USDC records a true $\mathrm{VaR}_{99\%}$ of 0.018032, which equates to a daily price decline of $e^{-0.018032} - 1 \approx 1.79\%$. The corresponding Expected Shortfall of 0.026408 indicates that on days when tail losses breach this threshold, the conditional average price contraction is bounded at $e^{-0.026408} - 1 \approx 2.61\%$. At the extreme five-year horizon, the $\mathrm{VaR}_{99.92\%}$ of 0.040118 maps to a potential price decline of $e^{-0.040118} - 1 \approx 3.93\%$, a result that demonstrates strong empirical consistency with the return levels reported in Table 5.12. Within this extreme multi-year tail scenario, if the VaR limit is violated, the conditional average loss is estimated by the ES (0.060648) to be a price decline of $e^{-0.060648} - 1 \approx 5.88\%$.

Lastly, USDT confirms the stability of the fiat peg, exhibiting minimal extreme tail risk. Its rescaled $\text{VaR}_{99\%}$ of 0.002179 corresponds to a negligible daily price decline of $e^{0.002179} - 1 \approx 0.22\%$, and even at the five-year horizon the $\text{VaR}_{99.92\%}$ of 0.007037 implies a maximum daily contraction of just $e^{0.007037} - 1 \approx 0.70\%$.

5.2.4 Model Evaluation

The evaluation of the GPD model fitted via the Peaks Over Threshold method is assessed using the same four standard diagnostic plots employed in the GEV analysis: the probability plot, the quantile plot, the return level plot, and the density plot [16]. These diagnostics are interpreted jointly to assess model adequacy. Overall, the POT method yields a more reliable characterization of tail risk than the Block Maxima approach, consistent with its more efficient use of extreme observations. However, the improvement is not uniform across assets, and in one instance, specifically the USDC loss series, the GPD fit is notably worse than its GEV counterpart. For brevity, only the most informative plots are presented here, while the full set of diagnostics is reported in Appendix B.

Bitcoin (BTC) The GPD model provides a good overall fit for Bitcoin in both the loss and gain series. The probability plot shows observations closely aligned with the diagonal, and the density plot confirms that the fitted distribution captures the empirical shape well. The most notable weakness is in the quantile plot, which exhibits some departures from the diagonal in the upper tail, indicating that the model tends to underestimate the most extreme observations. This pattern is consistent across both series, though it remains limited in magnitude and does not materially affect the overall adequacy of the fit. The return level plot is broadly consistent, with the majority of exceedances falling within the confidence bands. On balance, the POT approach offers an adequate representation of extreme behavior for this asset.

Ethereum (ETH) A similar conclusion holds for Ethereum. Both the probability and quantile plots display no clear evidence of systematic misspecification, and the model captures tail behavior reasonably well across both series. As observed in the GEV analysis, the gain series exhibits a marginally weaker fit than the loss series, with slightly more pronounced deviations in the upper tail of the quantile plot and a small number of exceedances falling outside the return level confidence bands. These discrepancies remain limited in magnitude and do not materially affect the overall adequacy of the model. The GPD thus provides a satisfactory characterization of extreme behavior for Ethereum, broadly in line with the GEV results for this asset.

Tether (USDT) For USDT the POT method yields results broadly similar to those of the GEV, which had already displayed a pronounced lack of fit for this asset. For the loss series, the probability plot is somewhat better aligned with the diagonal relative to the GEV case, suggesting a marginal improvement in the central fit. However, the quantile and return level plots continue to exhibit the same difficulties in tail estimation, with large outliers attributable to the ARMA–GARCH standardization remaining visible and breaching the confidence bands. A comparable pattern is observed for the gain series, shown in Figure 5.13, where the quantile and return level plots are largely indistinguishable from their GEV counterparts. On balance, no clear improvement over the GEV model can be established for USDT, suggesting that the structural characteristics of stablecoin returns, rather than the choice of EVT method, are the primary source of model inadequacy.

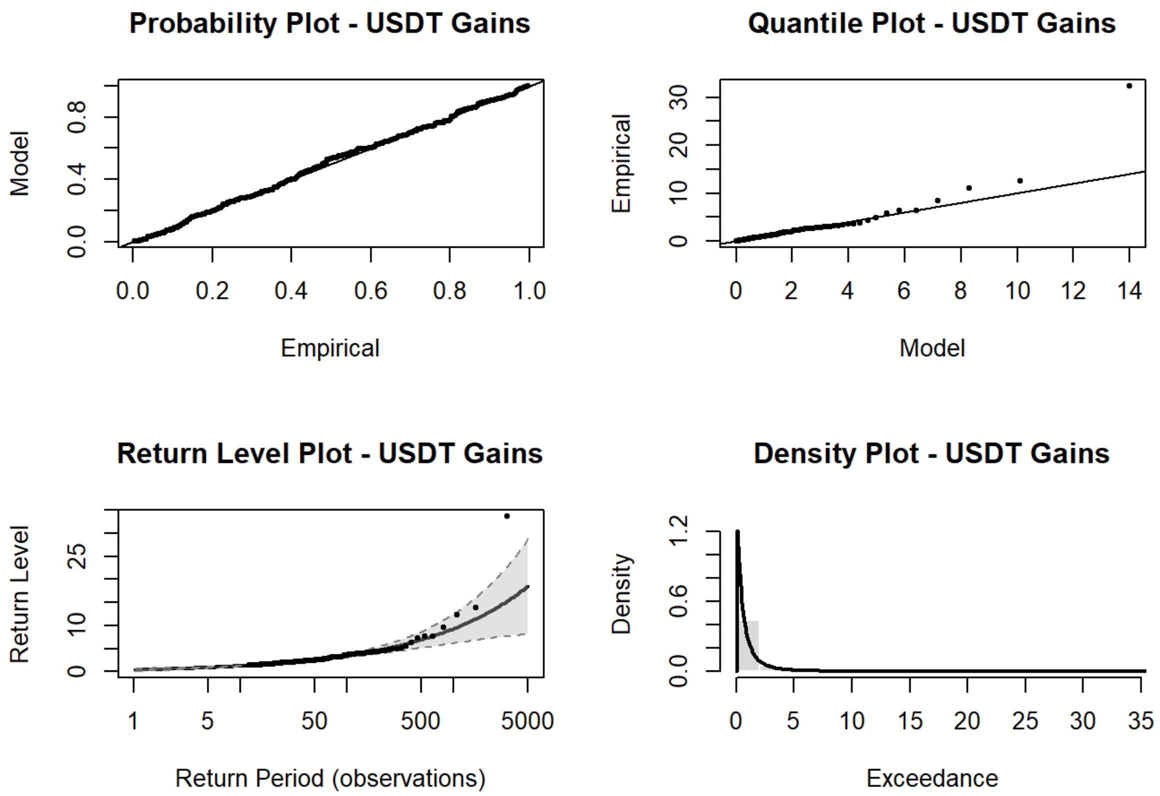

Figure 5.13: GPD diagnostic plots for USDT gains. As in the GEV case, large outliers attributable to the ARMA–GARCH standardization remain visible in the quantile and return level plots, breaching the 95% confidence band, while the probability plot shows only a marginal improvement in the central fit relative to the GEV model.

USD Coin (USDC) The results for USDC are mixed and present the most notable divergence between the two methods. For the gain series, the GPD model represents a clear improvement over the GEV fit: the probability and quantile plots show substantially better alignment with the diagonal, and while the return level plot still reveals some difficulty in estimating the extreme upper tail, the outliers are considerably less pronounced than those observed under the GEV framework. In contrast, the loss series tells a different story. The POT model for USDC losses performs markedly worse than its GEV counterpart, with clear deviations in both the probability and quantile plots and multiple exceedances falling outside the return level confidence bands. This is illustrated in Figure 5.14. This asymmetry between the two series is likely attributable to the structural nature of stablecoin returns combined with the limited number of observations available for USDC, which represents the shortest time series across all assets examined, reducing the reliability of tail estimation under both approaches.

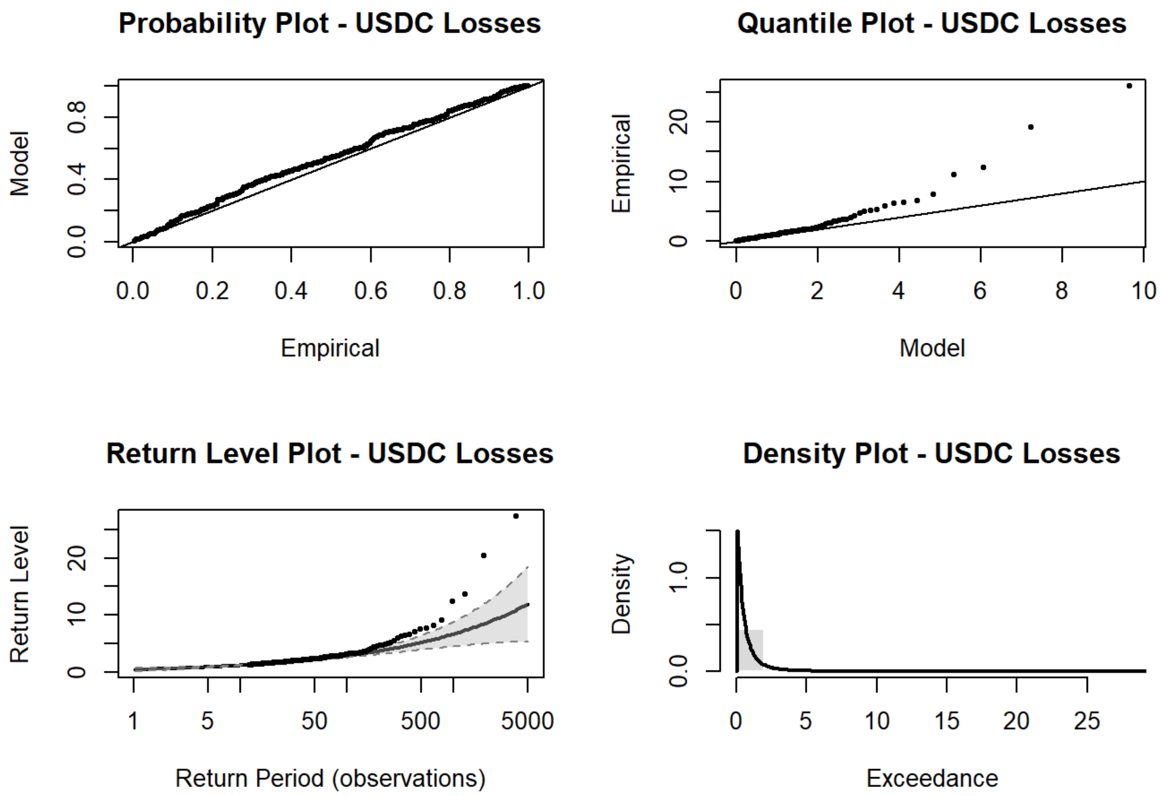

Figure 5.14: GPD diagnostic plots for USDC losses. In contrast to the gain series, where the POT model represents a clear improvement over the GEV, the loss series exhibits a markedly poorer fit, with substantial deviations in the probability and quantile plots and multiple exceedances falling outside the 95% confidence band in the return level plot.

Overall, the diagnostic analysis confirms a pattern broadly consistent with the GEV results. The GPD model provides an adequate fit for BTC and ETH, with the POT approach offering modest improvements over the Block Maxima method, consistent with its more efficient use of extreme observations. For the stablecoin series, however, neither method achieves a satisfactory fit. As discussed in the GEV analysis, stablecoins are inherently difficult to model with standard Extreme Value tools: their return distributions combine prolonged near-zero variation with sudden, short-lived spikes driven by de-pegging events, producing a tail structure that is highly atypical. Their comparatively shorter time series further compounds this challenge, reducing the precision of tail estimation regardless of the method employed. The one notable exception to the general pattern of improvement is the USDC loss series, where the GPD fit is unexpectedly worse than its GEV counterpart, highlighting that the advantages of the POT approach are not guaranteed and may be sensitive to the specific tail structure and threshold selection for individual series.

5.3 Final Remarks

The empirical results obtained from the Block Maxima (GEV) and Peaks-Over-Threshold (POT) methods show how different types of digital assets behave under extreme conditions. For both Tether (USDT) and USD Coin (USDC), the two statistical methods agree on two main points. First, the tail shape of these pegged assets belongs to the heavy-tailed Fréchet family ($\xi > 0$). Second, their risk measures (the return levels, VaR and ES) are very small, making them look almost invisible when compared to the high-risk profiles of Bitcoin and Ethereum. This accurately reflects the real-world economic mechanics of stablecoins. The mathematical framework successfully separates how severe a shock is compared to normal days from how much value the asset actually loses in dollars.

To understand this clearly, we must look at what the shape parameter ξ actually measures in Extreme Value Theory. The shape parameter is a number that does not tell us the actual size of a loss. Instead, it measures how unusual a shock is compared to the asset's normal daily price movements [16]. For Bitcoin and Ethereum, the analysis shows tails that are close to the Gumbel baseline or only slightly heavy. This happens because the two-step GARCH model appears to perform reasonably well for highly volatile assets. It captures and removes much of the continuous price changes that happen every day. Because this volatility filter cleans the data efficiently, the remaining variations show a smoother and more predictable tail behavior.

For stablecoins, this situation is completely reversed because of the artificial dollar peg. Assets like USDT and USDC experience very long periods of flat calm, during which the price barely moves and volatility is close to zero. This everyday flatness is only broken by rare, sudden crises when the stablecoin temporarily loses its peg to the dollar. When a de-pegging event happens, even a small price drop of 1% or 2% represents a massive shock compared to the long history of near-zero daily changes. The GARCH volatility filter divides these sudden jumps by an extremely low conditional variance, generating massive standardized innovations. Consequently, Extreme Value Theory identifies this huge disproportion in the filtered residuals, resulting in a positive ξ. The model is not saying that stablecoins are riskier than Bitcoin; it is simply pointing out that a de-pegging event is a major structural shock relative to the stablecoin's usual quiet behavior.

In conclusion, this analysis provides a valuable methodological proof of the limits and strengths of financial models in crypto markets. The combination of GARCH and EVT models is naturally designed to work best with continuous assets like Bitcoin and Ethereum, which move constantly and show a lot of volatility. Stablecoins break these standard assumptions because they alternate between absolute flatness and sudden jumps, which puts a lot of strain on standard variance models. However, both the GEV and POT frameworks focus on the tails of the distribution rather than on ordinary daily observations. While the results highlight several challenges in modeling assets with fundamentally different market dynamics, they also demonstrate the usefulness of EVT as a framework for investigating tail risk in cryptocurrency markets.

The following section extends the analysis to a multivariate framework, aiming to characterize the joint tail behavior of the four cryptocurrency series and assess whether extremes tend to occur concurrently across assets.

Chapter 6

Bivariate Extreme Value Theory Analysis

The bivariate Extreme Value analysis, as the univariate analysis, was conducted in `R` [59], using the `evd` package [63], which provides functions for fitting bivariate extreme-value dependence models. These include the Logistic, Bilogistic, and Dirichlet (Coles–Tawn) models, estimated via maximum likelihood with associated standard errors, as well as diagnostic and graphical tools for model assessment, such as χ and $\bar{\chi}$ plots.

6.1 Componentwise Maxima

This section presents the Componentwise Maxima analysis. As described in Chapter 3, estimation proceeds under two approaches: two-stage estimation and full likelihood estimation.

6.1.1 Two-Stage Estimation

The two-stage estimation begins by forming all pairwise combinations of the four cryptocurrencies, separately for the upper tail (gains) and the lower tail (losses), before applying the Fréchet transformation. Each pair of Block Maxima is then transformed to standard Fréchet margins following Equation 3.5, using the GEV parameter estimates $\hat{\mu}$, $\hat{\sigma}$, and $\hat{\xi}$ obtained in Section 5.1.1. The standardized data are then fitted with three dependence models: Logistic, Bilogistic, and Dirichlet [16].

Fitting is performed via the `fbvevd`[1] function in `R` [59]. Since `fbvevd` defaults to full likelihood estimation, jointly estimating both the marginal and dependence parameters, the two-stage approach requires the marginal parameters to be fixed at their standard Fréchet values ($\mu = 1$, $\sigma = 1$, $\xi = 1$), so that only the dependence parameter α (and β) is estimated in this second stage [16].

Gains Model selection is based on the Akaike Information Criterion (AIC) [3, 43]. Table 6.1 reports the AIC values for all three models across each asset pair for the gains.

The Bilogistic model failed to converge for the BTC–USDT and ETH–USDT pairs, likely because the near-independence between a stablecoin and a highly volatile asset leaves little dependence structure for an asymmetric model to capture.

The Logistic model achieves the lowest AIC across all asset pairs except BTC–USDC and ETH–USDC, where the Bilogistic yields a marginally lower AIC. For BTC–USDC the AIC

[1] https://www.rdocumentation.org/packages/evd/versions/2.3-7.1/topics/fbvevd

Table 6.1: AIC comparison of bivariate extreme-value dependence models. Two Step Approach (gains).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH–BTC	**1001.896**	1002.276	1003.983
BTC–USDT	**875.665**	—	878.762
BTC–USDC	719.615	**719.063**	719.833
ETH–USDT	**926.151**	—	927.319
ETH–USDC	741.272	**739.590**	740.939
USDC–USDT	**697.772**	699.765	699.579

Note: blue denotes the high-volatility asset (BTC/ETH); bold indicates the lowest AIC per pair.

difference is small ($\Delta\text{AIC} = 0.552$), and the standard error on the Bilogistic dependence parameter is degenerate ($\text{SE} \approx 0$). A similar behavior occurs under the Logistic specification for the ETH–USDC pair. Here, the estimate $\hat{\alpha} = 0.9991$ is practically equal to 1, which is the exact boundary of the model representing absolute independence. In Extreme Value Theory, when an estimate hits this boundary, the standard statistical assumptions break down. As a result, the formula used to calculate the standard error collapses, making it unreliable [16]. Given this lack of reliable inference, the small AIC improvement carries no real weight, and the Logistic model is retained across all pairs to keep the analysis consistent and easy to interpret.

Having selected the Logistic model uniformly, Table 6.2 reports the estimated dependence parameter $\hat{\alpha}$ and the corresponding extremal dependence coefficient for each pair. Recall that under the Logistic model $\alpha \in (0,1]$, with values near zero indicating strong dependence and values approaching one indicating near-independence; correspondingly, $\hat{\chi} = 2 - 2^{\alpha}$, where $\hat{\chi}$ close to zero signals near-independence [16].

Table 6.2: Parameter estimates and extremal dependence coefficients. Two Step Approach (gains).

Asset Pair	$\hat{\alpha}$	Std. Error	$\hat{\chi}$
ETH–BTC	0.5628	0.0408	0.5229
BTC–USDT	0.9121	0.0704	0.1182
BTC–USDC	0.9725	0.0762	0.0340
ETH–USDT	0.9342	0.0658	0.0892
ETH–USDC	0.9991	2×10^{-6}	0.0012
USDC–USDT	0.9489	0.0855	0.0695

The parametric estimates reveal a clear hierarchy in tail dependence. The ETH–BTC pair exhibits the strongest tail association ($\hat{\alpha} = 0.5628$, $\hat{\chi} = 0.5229$), validating the empirical co-movements observed between the two largest cryptocurrencies by market capitalization. To visually validate these parametric conclusions, the Block Maxima for gains transformed into standard Fréchet margins are plotted in the left panel of Figure 6.1. The plot provides a clear empirical justification of the estimates. For the ETH–BTC pair, a prominent alignment of observations along the main diagonal is visible, particularly for simultaneous large values. This concentration of co-occurring tail events strongly reflects the high estimated extremal dependence coefficient. On the other hand, observations near the origin represent periods where both assets experienced low maximum gains within the specific block, carrying no information regarding the loss tail, which is modeled independently.

The remaining pairs display some dependence under this framework, with $\hat{\alpha}$ values clustered between 0.91 and 0.97, apart from the ETH–USDC pair, which despite showing near-

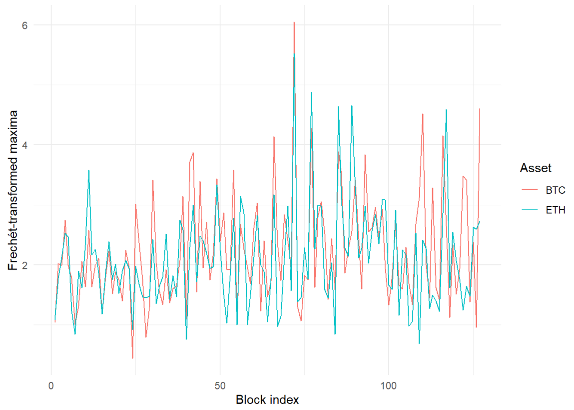

(a) ETH–BTC

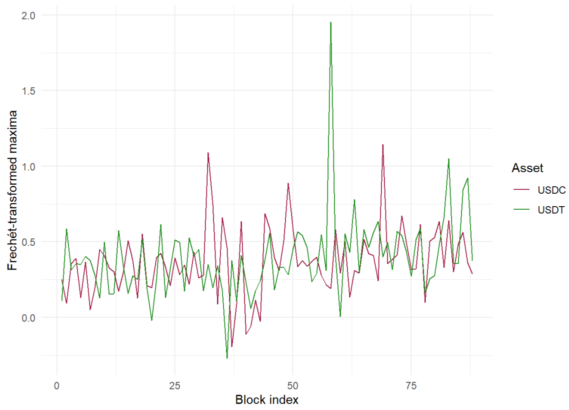

(b) USDC–USDT

Figure 6.1: Fréchet transformed Block Maxima for the ETH–BTC pair (left) and the USDC–USDT pair (right).

independence carries a degenerate standard error, making the estimate inconclusive. Dependence is comparatively stronger for the two pairs involving USDT and the two high-risk cryptocurrencies. The USDC–USDT stablecoin pair displays a weak tail dependence ($\hat{\alpha} = 0.9489$), suggesting that extreme deviations from the peg are not driven by a common shock. This empirical behavior is illustrated in the right panel of Figure 6.1, which presents the stablecoin joint extremes transformed into standard Fréchet margins.

Losses The same procedure is applied to the lower tail. Table 6.3 reports the AIC values for the three dependence models across each transformed Fréchet asset pair.

Table 6.3: AIC comparison of bivariate extreme-value dependence models. Two Step Approach (losses).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH–BTC	**1007.909**	1009.776	1009.041
BTC–USDT	**890.043**	890.262	891.369
BTC–USDC	736.089	735.896	**735.587**
ETH–USDT	**903.862**	906.312	906.216
ETH–USDC	**731.801**	733.651	—
USDC–USDT	**717.962**	718.398	718.355

Based on the AIC, the Logistic model is selected for almost all pairs, with the sole exception of BTC–USDC, for which the Dirichlet model (Coles-Tawn) yields a lower AIC. However, even for this pair, the Logistic model is retained. This decision is motivated by the fact that for both the Dirichlet and Bilogistic specifications, the estimated parameters hit the boundary of the parameter space, causing the standard error calculations to collapse. The final parametric estimates are then reported in Table 6.4.

The ETH–BTC pair exhibits a solid tail association ($\hat{\alpha} = 0.7012$, $\hat{\chi} = 0.3741$), though it is weaker than its upper-tail counterpart. This implies that extreme negative shocks are slightly less dependent than extreme positive shocks between the two main cryptocurrencies. The USDC–USDT stablecoin pair displays a moderate and significant tail dependence ($\hat{\alpha} = 0.7605$, $\hat{\chi} = 0.3058$), suggesting that extreme losses (de-pegging risks) tend to, intuitively, co-occur. While, the remaining volatile-stablecoin pairs display some weak tail association, with $\hat{\alpha}$ values ranging between 0.901 (BTC–USDT) and 0.973 (BTC–USDC). As before, the mixed pairs

Table 6.4: Parameter estimates and extremal dependence coefficients. Two Step Approach (losses).

Asset Pair	$\hat{\alpha}$	Std. Error	$\hat{\chi}$
ETH–BTC	0.7012	0.0483	0.3741
BTC–USDT	0.9006	0.0617	0.1331
BTC–USDC	0.9726	0.0824	0.0376
ETH–USDT	0.9394	0.0628	0.0822
ETH–USDC	0.9351	0.0718	0.0879
USDC–USDT	0.7605	0.0637	0.3058

involving USDT exhibit comparatively higher dependence; in this case, however, ETH–USDC also shows a non-negligible degree of dependence, with $\hat{\alpha} \approx 0.94$.

While this two-stage approach provides a valuable preliminary assessment, fixing the marginal distributions prior to estimating the dependence structure fails to account for the uncertainty in the marginal parameters, potentially leading to an understatement of the final standard errors. This limitation motivates the adoption of a joint full likelihood approach in the next section.

6.1.2 Full Likelihood Estimation

To compare against the two-stage procedure and avoid carrying forward the estimation errors of the univariate marginal fits into the dependence parameters, a full likelihood estimation framework is also implemented, where marginal GEV parameters and bivariate dependence parameters are optimized simultaneously.

Gains Applying the full likelihood approach to the gains, all three dependence models are fitted simultaneously across marginal and dependence parameters, as done before. Table 6.5 reports the resulting AIC values.

Table 6.5: AIC comparison of bivariate extreme-value dependence models. Full Likelihood Approach (gains).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH–BTC	**552.992**	553.580	555.217
BTC–USDT	**647.374**	649.180	649.260
BTC–USDC	**520.684**	532.476	520.907
ETH–USDT	**644.747**	645.395	645.829
ETH–USDC	508.323	**506.190**	507.605
USDC–USDT	**586.643**	—	588.627

Note: blue denotes the high-volatility asset (BTC/ETH); bold indicates the lowest AIC per pair.

Following the same pattern noted earlier, the Bilogistic specification for ETH–USDC produces a marginally lower AIC (ΔAIC ≈ 2.1) than the Logistic model, but both models yield a degenerate standard error on the dependence parameter, indicating that the joint tail observations are too few to reliably distinguish between the two specifications. All pairs are modeled here using the Logistic specification, matching the choice made throughout the rest of the analysis.

The full likelihood dependence parameters are provided in Table 6.6 provides full likelihood estimations that are broadly consistent with the two-stage approach, with only minor numerical

Table 6.6: Parameter estimates and extremal dependence coefficients. Full Likelihood Approach (gains).

Asset Pair	$\hat{\alpha}$	Std. Error	$\hat{\chi}$
ETH–BTC	0.5501	0.0486	0.5359
BTC–USDT	0.9153	0.0718	0.1141
BTC–USDC	0.9825	0.0763	0.0242
ETH–USDT	0.9272	0.0706	0.0988
ETH–USDC	0.9995	2×10^{-6}	0.0007
USDC–USDT	0.9682	0.0831	0.0436

differences in the dependence structure. The ETH–BTC pair remains the most strongly associated ($\hat{\alpha} = 0.5501$, $\hat{\chi} = 0.5359$), with both the point estimate and the implied dependence coefficient closely matching those obtained in Section 6.1.1. All remaining pairs continue to cluster near asymptotic independence, with $\hat{\alpha} > 0.91$ for every volatile–stablecoin combination. This is economically sensible: since stablecoins are pegged near unity, large positive returns in BTC or ETH need not be accompanied by comparably large positive returns in USDT or USDC. The ETH–USDC pair again carries a degenerate standard error, and so the $\hat{\alpha}$ estimate is not a confirmation of near-independence.

Losses The joint likelihood maximization routine is replicated for the lower tail. Table 6.7 and Table 6.8 report the model selection criterion and parameter estimates, respectively.

Table 6.7: AIC comparison of bivariate extreme-value dependence models. Full Likelihood Approach (losses).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH–BTC	**607.313**	609.302	608.464
BTC–USDT	**674.340**	681.742	676.280
BTC–USDC	505.165	**503.803**	504.157
ETH–USDT	**660.544**	663.366	662.527
ETH–USDC	**486.897**	492.334	488.918
USDC–USDT	**567.321**	567.841	567.828

The Logistic model has a lower AIC for all pairs except BTC–USDC, where the Bilogistic yields a marginally lower AIC (ΔAIC ≈ 1.4). As in the gains case for ETH–USDC, inspection of the Bilogistic output for BTC–USDC reveals a degenerate standard error on the dependence parameter, indicating numerical instability at the boundary of the parameter space. The small AIC improvement therefore carries no reliable inferential weight, and the Logistic model is retained. Notably, the ETH–USDC pair, which was a boundary case in the gains sample, behaves regularly here, with the Logistic model selected and well-behaved standard errors. The Logistic specification is therefore maintained across all pairs for consistency.

The negative tail parameters confirm the overall structure observed in the upper tail, with some notable differences. ETH–BTC remains the most dependent pair ($\hat{\alpha} = 0.6915$, $\hat{\chi} = 0.3849$), though dependence is weaker than in the gains sample, suggesting that extreme losses in the two assets are less synchronized than extreme gains. The most interesting reversal concerns USDC–USDT, which shows stronger lower-tail dependence ($\hat{\alpha} = 0.7649$, $\hat{\chi} = 0.3011$) than in the upper tail. This means that during periods of systemic stress, a loss of confidence in

Table 6.8: Parameter estimates and extremal dependence coefficients. Full Likelihood Approach (losses).

Asset Pair	$\hat{\alpha}$	**Std. Error**	$\hat{\chi}$
ETH–BTC	0.6915	0.0556	0.3849
BTC–USDT	0.9078	0.0609	0.1238
BTC–USDC	0.9962	0.0905	0.0053
ETH–USDT	0.9360	0.0656	0.0875
ETH–USDC	0.9223	0.0754	0.1051
USDC–USDT	0.7649	0.0721	0.3011

fiat-pegged instruments tends to affect both stablecoins simultaneously, producing correlated extreme losses that are absent during normal or positive market conditions.

Overall, the bivariate Componentwise Maxima analysis delivers a consistent and interpretable picture. The most robust finding across both estimation methods and both tails is the dependence between assets of similar volatility profile: the ETH–BTC pair exhibits meaningful tail dependence in both the gains and losses samples under both the two-stage and full likelihood approaches, confirming that extreme movements in the two major cryptocurrencies tend to occur together. Moreover, dependence is consistently higher in the gains than in the losses under both estimation methods, suggesting that joint extreme upward movements are more strongly linked than joint extreme downward movements. The USDC–USDT pair follows a related pattern in the lower tail, where systemic stress appears to drive correlated extreme losses between the two stablecoins.

For the volatile–stablecoin pairs, the evidence is less conclusive: a weak but coherent degree of dependence emerges for the mixed pairs involving USDT, namely BTC–USDT and ETH–USDT, consistent across both the gains and losses samples. This pattern is also reflected in the full likelihood estimation, but in this case for ETH–USDC, where some dependence persists in the losses. For the remaining pairs, however, no significant dependence is detected. The main takeaway is therefore that *same-volatility* asset pairs display detectable tail dependence, while *cross-volatility* pairs provide only limited and inconsistent evidence of systematic joint extreme behavior.

This naturally leads to an investigation of whether volatile and stable assets exhibit *opposing* tail behavior: specifically, whether extreme losses in BTC or ETH coincide with extreme gains in USDT or USDC, consistent with investors rotating into stable assets during periods of market stress, the so-called safe haven hypothesis. The following part of the chapter addresses this directly by extending the analysis to the gain and loss mixed-tail setting.

6.1.3 Asymmetric Tail Analysis

Given that full likelihood estimation provides asymptotically unbiased and statistically efficient estimates, this framework is extended to analyze mixed-tail dynamics.A core objective of this analysis is to investigate asset reallocation dynamics during periods of extreme market distress. Specifically, we test the safe haven hypothesis: whether extreme losses in volatile cryptocurrencies systematically trigger simultaneous positive capital inflows into stablecoins, causing positive tail deviations in the latter due to short-term peg inelasticity. To implement this, bivariate models are fitted by pairing the negative Block Maxima (losses) of BTC/ETH with the positive Block Maxima (gains) of USDT/USDC, following the reasoning outlined in Section 3.4.

Table 6.9 details the cross-tail parametric model selection via AIC.

Table 6.9: AIC comparison of bivariate extreme-value dependence models. Full Likelihood Approach (mixed tails).

Asset Pair	Logistic	Bilogistic	Dirichlet
BTC–USDT	665.8553	666.0987	**665.7946**
BTC–USDC	533.6655	**528.6823**	529.9619
ETH–USDT	**652.3510**	653.3552	653.4324
ETH–USDC	**519.4199**	—	521.4193

Note: blue denotes the high-volatility asset (BTC/ETH); bold indicates the lowest AIC per pair.

The Logistic model is selected for ETH–USDT and ETH–USDC. For BTC–USDT the Dirichlet yields a negligible improvement (ΔAIC ≈ 0.06), and for BTC–USDC the Bilogistic yields a larger but still moderate improvement (ΔAIC ≈ 4.98). However, inspection of the Bilogistic output for BTC–USDC reveals again a degenerate standard error on the α parameter, indicating numerical instability at the boundary of the parameter space and rendering the asymmetric specification unreliable for this pair. As in the previous cases, the AIC improvement carries no reliable inferential weight under these conditions. We therefore fit the Logistic model uniformly across all four pairs, so that any differences in results reflect the data rather than the choice of specification.

Table 6.10: Parameter estimates and extremal dependence coefficients. Full Likelihood Approach (mixed tails).

Asset Pair	$\hat{\alpha}$	Std. Error	$\hat{\chi}$
BTC–USDT	0.8924	0.0660	0.1437
BTC–USDC	0.9222	0.0611	0.1050
ETH–USDT	0.9419	0.0712	0.0790
ETH–USDC	0.9692	0.0686	0.0423

The BTC–USDT pair shows the strongest cross-tail dependence ($\hat{\alpha} = 0.8924$, $\hat{\chi} = 0.1437$), and BTC–USDC follows a similar pattern ($\hat{\chi} = 0.1050$). Both values exceed the corresponding same-tail coefficients for these pairs, where gains-gains and losses-losses structures converged near independence. This suggests that the cross-tail relationship, extreme BTC losses coinciding with extreme stablecoin gains, is more pronounced than any same-direction co-movement, which is precisely what the safe haven hypothesis predicts. The ETH–stablecoin pairs show weaker but non-negligible cross-tail dependence: while $\hat{\chi} = 0.0790$ for ETH–USDT and $\hat{\chi} = 0.0423$ for ETH–USDC are small, they remain above the near-zero same-tail estimates for these pairs, suggesting a faint directional signal that falls short of constituting evidence for the safe haven hypothesis. Overall, the evidence points to a Bitcoin-specific dynamic, whereby extreme BTC losses are accompanied by extreme stablecoin gains in roughly 10% to 14% of tail episodes, while the ETH signal, though not absent, is too weak to draw firm conclusions.

To complement these results and provide a comprehensive analysis of the joint tails, we mirror our previous univariate strategy and transition to the threshold approach, completing the bivariate framework.

6.2 Threshold Approach

The threshold approach is carried out following the same two estimation strategies as the Componentwise Maxima analysis: two-stage and full likelihood estimation.

6.2.1 Two-Stage Estimation

The rationale behind the two-stage approach mirrors the one described for the Block Maxima framework, following the theory presented in Chapter 3. Each exceedance of each cryptocurrency series is transformed to standard Fréchet margins via Equations 3.9 and 3.10, using the GPD parameter estimates $\hat{\sigma}$ and $\hat{\xi}$ obtained in Chapter 5. The standardized exceedances are then paired for both gains and losses to analyze the bivariate dependence structure.

The transformed data are fitted via the Logistic, Bilogistic, and Dirichlet models [16], with model evaluation performed through the AIC criterion. Fitting is performed via the `fbvpot`[2] function in `R` [59], with marginal parameters fixed at their standard Fréchet values.

Gains As shown in Table 6.11, the Bilogistic and Dirichlet models suffer from a systematic failure to converge across almost all asset pairs.

Table 6.11: AIC comparison of bivariate extreme-value dependence models. Two Step Approach (gains).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH – BTC	**16910.11**	16910.79	16943.40
BTC – USDT	**13096.91**	—	—
BTC – USDC	**10008.81**	—	—
ETH – USDT	**13682.31**	—	—
ETH – USDC	**10465.92**	—	—
USDC – USDT	**8148.90**	—	—

Note: blue denotes the high-volatility asset (BTC/ETH); bold indicates the lowest AIC per pair.

This numerical breakdown occurs because these asymmetric frameworks collapse when applied to assets that exhibit structural asymptotic independence, which is typically the case for pairs involving stablecoins (USDT and USDC). These models are only able to achieve convergence for the ETH–BTC pair, which displays a strong and genuine mutual tail dependence. However, even for the ETH–BTC combination, the classical Logistic model is strictly preferred according to the AIC criterion, and a symmetric tail dependence structure is therefore assumed for all analyzed pairs [45].

The corresponding Logistic parameter estimates are given in Table 6.12.

The ETH–BTC pair again emerges as the most dependent combination ($\hat{\alpha} = 0.7402$, $\hat{\chi} = 0.3296$), indicating that extreme gains in the two major cryptocurrencies tend to co-occur. For all remaining pairs, however, the estimated $\hat{\alpha}$ values cluster close to unity and the corresponding $\hat{\chi}$ coefficients are negligible, signaling near-independence. Notably, the threshold approach does not alter the qualitative conclusions obtained from the Componentwise Maxima analysis. Apart from the ETH–BTC pair, all remaining combinations continue to exhibit only weak extremal dependence, with the POT estimates lying even closer to the independence boundary.

[2] https://www.rdocumentation.org/packages/evd/versions/2.3-7.1/topics/fbvpot

Table 6.12: Parameter estimates and extremal dependence coefficients. Two Step Approach (gains).

Asset Pair	$\hat{\alpha}$	**Std. Error**	$\hat{\chi}$
ETH–BTC	0.7402	0.0101	0.3296
BTC–USDT	0.9871	0.0036	0.0178
BTC–USDC	0.9946	0.0033	0.0075
ETH–USDT	0.9814	0.0042	0.0256
ETH–USDC	0.9975	0.0028	0.0035
USDC–USDT	0.9867	0.0042	0.0184

Losses The same convergence pattern is observed for the lower tail: the Bilogistic and Dirichlet models fail to converge across all pairs, and the Logistic model is retained throughout. Table 6.13 reports the parameter estimates.

Table 6.13: Parameter estimates and extremal dependence coefficients. Two Step Approach (losses).

Asset Pair	$\hat{\alpha}$	**Std. Error**	$\hat{\chi}$
ETH–BTC	0.8381	0.0084	0.2123
BTC–USDT	0.9802	0.0044	0.0272
BTC–USDC	0.9852	0.0043	0.0204
ETH–USDT	0.9857	0.0038	0.0197
ETH–USDC	0.9830	0.0046	0.0234
USDC–USDT	0.9854	0.0041	0.0201

The ETH–BTC pair remains the only combination displaying meaningful tail dependence ($\hat{\alpha} = 0.8381$, $\hat{\chi} = 0.2123$), though its estimated dependence is weaker than in the upper tail, consistent with the findings from the Componentwise Maxima analysis. All remaining pairs exhibit near-complete independence, with $\hat{\chi}$ values collapsing to approximately 0.02 across the board. As in the gains case, the estimates obtained via the threshold approach are generally smaller than their Componentwise Maxima counterparts, particularly for pairs involving stablecoins. However, both approaches lead to the same qualitative conclusion: tail dependence remains concentrated in the ETH–BTC pair and is stronger for gains than for losses, while the remaining combinations lie close to asymptotic independence.

6.2.2 Full Likelihood Estimation

To address the limitations of the two-stage procedure, a full likelihood estimation framework is implemented, jointly optimizing the GPD marginal parameters and the bivariate dependence parameters.

Gains Under full likelihood estimation, the numerical difficulties encountered in the two-stage approach are noticeably mitigated, and the Bilogistic and Dirichlet models achieve convergence for a larger number of pairs. Table 6.14 reports the resulting AIC values.

The Logistic model is selected for the majority of pairs, apart from ETH–BTC and BTC–USDC, where the Bilogistic yields a lower AIC ($\Delta\text{AIC} \approx 1.7$ and $\Delta\text{AIC} \approx 26.7$, respectively), suggesting the presence of asymmetric tail dependence for these two combinations. We recall that the selection of an asymmetric model indicates that the tail dependence of Bitcoin

Table 6.14: AIC comparison of bivariate extreme-value dependence models. Full Likelihood Approach (gains).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH–BTC	5049.004	**5047.308**	5060.970
BTC–USDT	**5929.931**	5931.414	5931.416
BTC–USDC	4746.414	**4719.744**	4720.112
ETH–USDT	**5771.061**	5772.629	—
ETH–USDC	**4576.579**	—	—
USDC–USDT	**5250.596**	5254.553	5279.020

Note: blue denotes the high-volatility asset (BTC/ETH); bold indicates the lowest AIC per pair.

on Ethereum is not identical to that of Ethereum on Bitcoin. In the Bilogistic framework, the difference between the two estimated parameters measures the strength and direction of this asymmetry, capturing how one asset exerts a stronger influence on the joint tail than the other. However, inspection of the BTC–USDC output reveals degenerate standard errors on both the Logistic and Bilogistic dependence parameters (SE ≈ 0), signaling that the joint tail observations are too scarce to reliably distinguish between the two specifications and that the optimizer is being pushed toward the boundary of the parameter space [16]. Given this numerical instability, the Logistic model is kept for BTC–USDC to preserve consistency across the analysis. The full likelihood parameter estimates are reported in Table 6.15.

Table 6.15: Parameter estimates and extremal dependence coefficients. Full Likelihood Approach (gains).

Asset Pair	Model	Parameter	Estimate	Std. Error	$\hat{\chi}$
ETH–BTC	Bilogistic	α	0.5737	0.0295	0.5576
		β	0.4754	0.0335	
BTC–USDT	Logistic	α	0.9836	0.0115	0.0226
BTC–USDC	Logistic	α	0.9993	2.002×10^{-6}	0.0010
ETH–USDT	Logistic	α	0.9869	0.0126	0.0180
ETH–USDC	Logistic	α	0.9992	2.002×10^{-6}	0.0011
USDC–USDT	Logistic	α	0.9985	2.003×10^{-6}	0.0021

Genuine tail dependence under extreme positive shocks is restricted exclusively to the ETH–BTC pair. Under the Bilogistic framework, both asymmetry parameters are significantly bounded away from unity ($\hat{\alpha} = 0.5737$, $\hat{\beta} = 0.4754$), yielding an extremal dependence coefficient of $\hat{\chi} = 0.5576$. From an economic perspective, this implies that when either Bitcoin or Ethereum experiences an extreme positive return shock, there is approximately a 55.8% probability that the other major asset will simultaneously experience a tail event of comparable magnitude. The symmetric Logistic specification, retained for comparability, corroborates this finding with an estimated $\hat{\alpha} = 0.5278$, confirming substantial dependence independently of the asymmetric parameterization.

All pairs involving a stablecoin exhibit a marked weakening of tail dependence, though the strength of this conclusion differs across pairs. For BTC–USDT and ETH–USDT, the Logistic dependence parameter $\hat{\alpha}$ converges to values near its upper theoretical limit (0.9836 and 0.9869, respectively) with well-behaved standard errors, and the corresponding $\hat{\chi}$ coefficients collapse to near-zero levels (0.0226 and 0.0180), supporting a reasonably reliable conclusion of asymptotic independence for these two pairs. For the pairs involving USDC (BTC–USDC, ETH–USDC, and USDC–USDT) $\hat{\alpha}$ similarly converges close to 1, but the associated standard errors are degenerate for all three. As discussed above, this reflects the breakdown of the normal approximation at the boundary of the parameter space rather than genuine estimation precision, so the apparent near-independence for these USDC pairs should be read as inconclusive.

Losses Table 6.16 reports the AIC values for the lower tail. The Logistic model is selected for most pairs, with the Bilogistic preferred for ETH–BTC ($\Delta\text{AIC} \approx 6.8$) and BTC–USDC ($\Delta\text{AIC} \approx 12.3$), and the Dirichlet preferred for ETH–USDC ($\Delta\text{AIC} \approx 3.3$).

Table 6.16: AIC comparison of bivariate extreme-value dependence models. Full Likelihood Approach (gains).

Asset Pair	Logistic	Bilogistic	Dirichlet
ETH–BTC	5895.513	**5888.703**	5891.422
BTC–USDT	**6031.153**	6033.151	6045.345
BTC–USDC	4524.551	**4512.219**	4513.166
ETH–USDT	**6010.192**	6013.278	—
ETH–USDC	4513.230	—	**4509.916**
USDC–USDT	**5116.590**	5117.891	5117.888

As in the gains case, the BTC–USDC and ETH–USDC pairs require careful interpretation. For BTC–USDC, both the Bilogistic and Logistic specifications produce degenerate standard errors ($\text{SE} \approx 0$) on the dependence parameters. The asymmetric specification is therefore rejected in favor of the Logistic model. For ETH–USDC, by contrast, the Logistic model selected in the losses sample is well-behaved, with regular standard errors, in contrast to its boundary behavior in the gains sample; the Logistic specification is retained here as well, for consistency. The complete parameter estimates are reported in Table 6.17.

Table 6.17: Parameter estimates and extremal dependence coefficients. Full Likelihood Approach (gains).

Asset Pair	Model	Parameter	Estimate	Std. Error	$\hat{\chi}$
ETH–BTC	Bilogistic	α	0.5827	0.0460	0.3855
		β	0.7621	0.0263	
BTC–USDT	Logistic	α	0.9842	0.0142	0.0218
BTC–USDC	Logistic	α	0.9996	2.002×10^{-6}	0.0006
ETH–USDT	Logistic	α	0.9953	0.0119	0.0066
ETH–USDC	Logistic	α	0.98200	0.01281	0.0248
USDC–USDT	Logistic	α	0.9852	0.0119	0.0204

The loss tail results confirm the structure observed in the gains analysis, with one notable difference. The ETH–BTC pair again exhibits robust tail dependence under the Bilogistic framework ($\hat{\alpha} = 0.5827$, $\hat{\beta} = 0.7621$, $\hat{\chi} = 0.3855$), and the symmetric Logistic estimate ($\hat{\alpha} = 0.6903$) corroborates this finding. The notable difference with respect to the gains sample concerns the asymmetry of the Bilogistic parameterization: the substantially larger $\hat{\beta}$ relative to $\hat{\alpha}$ suggests that extreme losses in ETH are more likely to be accompanied by extreme losses in BTC than vice versa, a directional structure absent from the upper tail.

All pairs involving a stablecoin are almost independent, apart from BTC–USDC, which again carries a degenerate standard error and therefore yields no conclusive information, with $\hat{\alpha}$ values ranging from 0.9842 to 0.9996 and $\hat{\chi}$ coefficients collapsing to near-zero levels. From an econometric standpoint, extreme crashes in Bitcoin or Ethereum do not systematically trigger simultaneous tail movements in stablecoins, whose return variation is bounded by their pegging mechanism.

Visual Independence

For the threshold approach, we can also display the empirical plots of χ and $\bar{\chi}$ to assess the tail dependence structure visually.

In the empirical analysis, the non-parametric estimation of these measures is computed via the `chiplot` [3] routine. To effectively interpret the resulting diagnostic output, the graphical components are defined as follows:

- *The X-axis (Threshold u)*: Represents the scale of shock severity, bounded between 0 and 1. As the threshold moves toward the right edge of the plot ($u \to 1$), the analysis discards ordinary fluctuations to focus exclusively on periods of severe market stress.
- *The Left Y-axis (χ plot):* Displays the empirical strength of the observed conditional tail dependence. A decaying pattern where the curve approaches zero as $u \to 1$ provides direct evidence of asymptotic independence.
- *The Right Y-axis ($\bar{\chi}$ plot):* Identifies the exact extremal dependence regime. If the empirical curve stabilizes strictly below the horizontal dashed line at 1, it mathematically confirms the presence of asymptotic independence . In this case, the exact height of the curve on the Y-axis quantifies the strength of the residual tail dependence.
- *The Dashed Control Lines:* Represent the corresponding confidence intervals (95% level). If the empirical curve falls within the bounds encompassing zero, it indicates that there is statistically no evidence of any relationship between the bivariate tails.

Figure 6.2 and Figure 6.3 display the empirical estimates of χ and $\bar{\chi}$ for the ETH–BTC and USDC–USDT pairs across both gains and losses.

[3] `https://www.rdocumentation.org/packages/evd/versions/2.3-7.1/topics/chiplot`

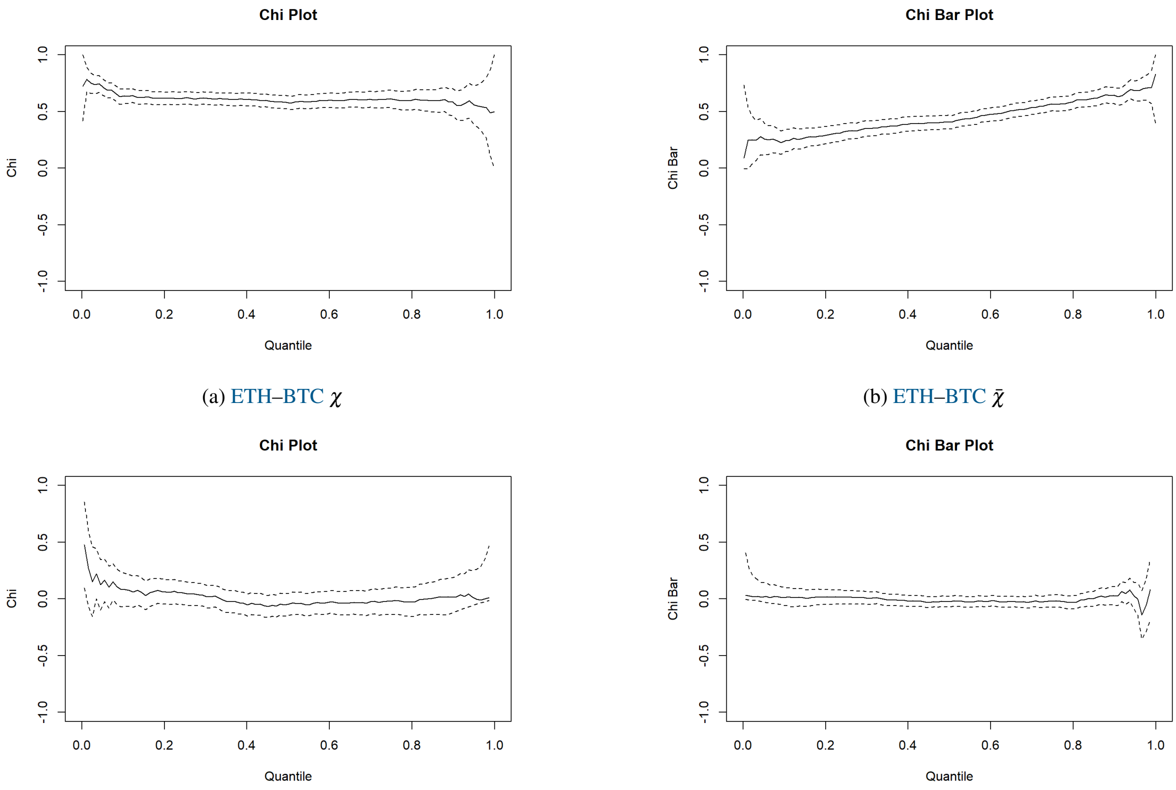

(a) ETH–BTC χ

(b) ETH–BTC $\bar{\chi}$

(c) USDT–USDC χ

(d) USDT–USDC $\bar{\chi}$

Figure 6.2: χ plot and $\bar{\chi}$ plot of the two homogeneous pairs (gains).

These specific pairs are selected to illustrate the two opposite regimes of extreme co-movements within homogeneous asset classes: ETH–BTC represents the high-volatility cryptocurrency cluster, while USDC–USDT represents the fiat-pegged stablecoin compartment. For brevity, the plots for all remaining asset combinations are provided in Appendix C.

For the ETH–BTC pair in the gains sample, Figure 6.2,the χ curve stabilizes at approximately 0.6 as $u \to 1$, and the $\bar{\chi}$ curve converges toward unity, jointly confirming asymptotic dependence with an empirical tail association of magnitude 0.6. For the USDC–USDT pair, both curves are consistent with asymptotic independence across the entire threshold range, corroborating the parametric estimates.

The losses plots confirm the same qualitative picture. The ETH–BTC pair displays residual tail dependence in both the χ and $\bar{\chi}$ panels, while the USDC–USDT pair shows no significant evidence of joint tail behavior across either measure.

Taken together, the threshold framework supports the main conclusion of the Componentwise Maxima analysis. Significant tail dependence is robustly detected for the ETH–BTC pair under both estimation approaches and in both tails. For the USDC–USDT pair, which the full likelihood GEV analysis tentatively identified as mildly dependent in the lower tail (Section 6.1.2), the threshold framework provides no supporting evidence: neither the parametric estimates nor the non-parametric χ and $\bar{\chi}$ diagnostics allow the hypothesis of independence to be discarded. The pegging mechanism therefore appears to effectively anchor the two stablecoins in the extremes, consistently with the univariate findings of Chapter 5.

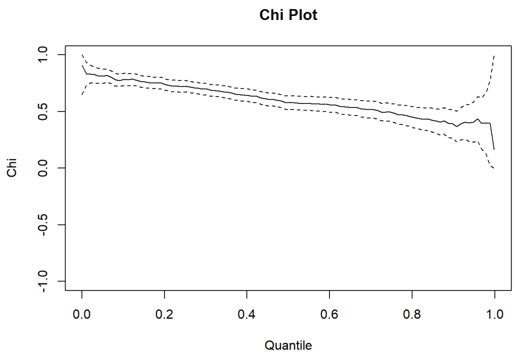

(a) ETH–BTC χ

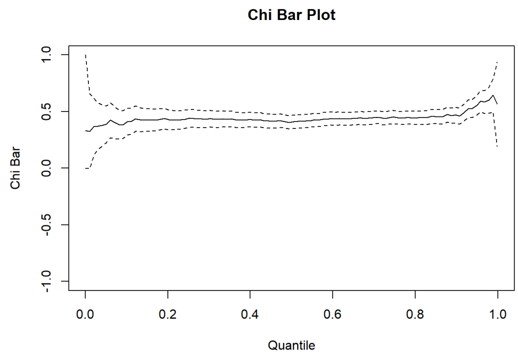

(b) ETH–BTC $\bar{\chi}$

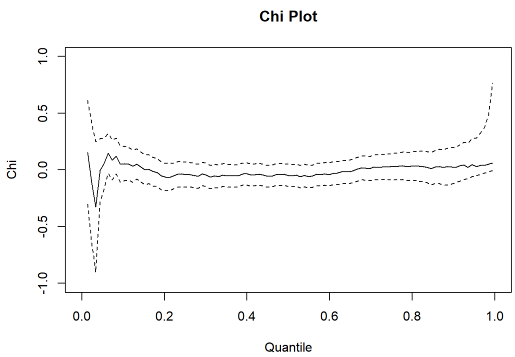

(c) USDT–USDC χ

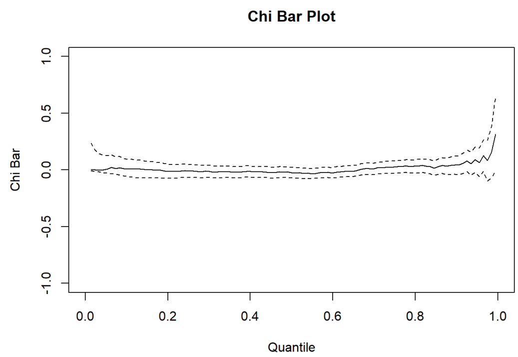

(d) USDT–USDC $\bar{\chi}$

Figure 6.3: χ plot and $\bar{\chi}$ plot of the two homogeneous pairs (losses).

6.2.3 Asymmetric Tails Analysis

The cross-tail analysis is replicated within the threshold framework, pairing the loss exceedances of BTC and ETH with the gain exceedances of USDT and USDC under full likelihood estimation.

Table 6.18 reports the AIC values for the three dependence models across all four pairs.

Table 6.18: AIC comparison of bivariate extreme-value dependence models. Full Likelihood Approach (mixed tails).

Asset Pair	**Logistic**	**Bilogistic**	**Dirichlet**
BTC–USDT	**6024.830**	6025.612	6025.621
BTC–USDC	**4709.460**	4711.460	4711.536
ETH–USDT	**6000.696**	6002.409	6005.212
ETH–USDC	**4707.510**	4709.216	4709.258

Note: blue denotes the high-volatility asset (BTC/ETH); bold indicates the lowest AIC per pair.

The Logistic model is selected for all four pairs, having lower AIC than the asymmetric alternatives. The parameter estimates are reported in Table 6.19.

All four $\hat{\alpha}$ estimates approach unity and the corresponding $\hat{\chi}$ coefficients are negligible, so the threshold framework finds no formal evidence of cross-tail dependence between extreme losses in volatile cryptocurrencies and extreme gains in stablecoins. This contrasts with the GEV mixed tails results, where the BTC–stablecoin pairs yielded $\hat{\chi}$ values of 0.14 and 0.11, sufficient to support a partial safe haven hypothesis.

Table 6.19: Parameter estimates and extremal dependence coefficients. Full Likelihood Approach (mixed tails).

Asset Pair	$\hat{\alpha}$	Std. Error	$\hat{\chi}$
BTC–USDT	0.9855	0.0113	0.01998
BTC–USDC	0.9524	0.0150	0.06482
ETH–USDT	0.9825	0.0110	0.02413
ETH–USDC	0.9616	0.0152	0.05257

Nevertheless, the cross-tail estimates are not entirely uninformative. The $\hat{\chi}$ values for BTC–USDC and ETH–USDC (0.0648 and 0.0526, respectively) are notably above the near-zero baseline recorded for the same-tail stablecoin pairs in the threshold framework, where $\hat{\chi} \approx 0.02$ across the board. The pattern is therefore at least consistent with the safe haven hypothesis, even if the threshold framework cannot confirm it at a statistically meaningful level.

6.3 Final Remarks

Taken together, both the Componentwise Maxima and the threshold approach consistently identify significant tail dependence between ETH and BTC, with the upper tail displaying stronger co-movement than the lower tail across both frameworks. The Componentwise Maxima analysis additionally detects marginal lower-tail dependence for the USDC–USDT pair, a signal that does not survive the more demanding threshold test. For all volatile–stablecoin pairs, both frameworks point toward asymptotic independence in the same-tail setting, broadly consistent with the effectiveness of the pegging mechanism under extreme market conditions. This conclusion, however, must be qualified for the threshold framework specifically. Under full likelihood estimation, several of the pairs involving a stablecoin return a dependence parameter pinned at the boundary of the parameter space alongside a degenerate standard error. This pattern is itself structurally informative: it arises because joint threshold exceedances, observations in which both series exceed their respective thresholds simultaneously, are extremely rare for pairs involving a stablecoin, a direct consequence of the pegging mechanism bounding the frequency and magnitude of extreme deviations.

Regarding the safe haven hypothesis, the evidence is suggestive but not conclusive. The Componentwise Maxima mixed tails analysis yields a partial signal for the BTC–stablecoin pairs, while the threshold framework does not confirm it at a statistically meaningful level. That said, the cross-tail $\hat{\chi}$ estimates remain systematically above the same-tail stablecoin baseline in both frameworks, and the directionality is consistent with the hypothesis. This residual signal, however modest, is noteworthy given the structural constraints of the data: stablecoins exhibit extremely limited return variation by design, and genuine joint tail episodes are correspondingly rare. In this context, the fact that cross-tail pairs display higher dependence than same-tail stablecoin pairs, even marginally, constitutes a result worth recording.

Chapter 7

Discussion and Limitations

The empirical findings presented in Chapters 5 and 6 demonstrate the utility of Extreme Value Theory (EVT) in characterizing tail risk dynamics and co-movements within the cryptocurrency ecosystem. However, several methodological and structural constraints must be addressed to contextualize these results properly and guide future extensions.

The first and most significant limitation concerns the GARCH filtering stage. The two-step procedure applied rests on the assumption that the filtered standardized residuals are approximately i.i.d. [51], a condition that is satisfied reasonably well for Bitcoin and Ethereum but is substantially harder to meet for stablecoins. Stablecoin historical data contains an embedded regime change: the early years of both USDT and USDC featured wide and frequent deviations from the dollar parity, while the mature period exhibits a near-flatline price series with volatility close to zero. A standard GARCH model fitted to the full sample period is forced to average across these two regimes, producing parameter estimates that are biased upward by the historical volatility of the early period [36]. The consequence is a conditional variance that is systematically overestimated in the modern period, causing contemporary standardized residuals to be deflated relative to their true scale. When a genuine de-pegging episode occurs, its magnitude is inconsistent with the tail behavior implied by the GEV-GPD parameters estimated from the mature low-volatility regime. This leads to extreme observations that dominate the tail diagnostics and emerge as pronounced macro-outliers. Future research could address this by replacing the standard GARCH specification with variants that explicitly model structural regime changes, such as Markov-Switching GARCH or Change-Point GARCH formulations, which would allow the baseline variance to step down following the maturation of the peg without contaminating the tail estimates with historical instability [5]. Consequently, the use of a more accurate, specifically tailored GARCH specification would have mitigated the computational bias in the estimation of the shape parameter (ξ). By avoiding the artificial compression of standardized residuals, the model would not have been forced to interpret individual depegging episodes as isolated macro-outliers, thereby reducing the tendency of the Extreme Value Theory (EVT) algorithm to incorrectly classify stablecoins within a heavy-tailed Fréchet distribution ($\xi > 0$).

The second major limitation arises from substantial differences in the historical lifespans of the selected assets, which result in a pronounced data imbalance across the sample. While Bitcoin and Ethereum offer long, rich time series spanning over a decade of market development, stablecoins are considerably younger. For instance, USD Coin (USDC) was launched in late 2018, restricting the synchronized empirical analysis to a compressed time window compared to the broader history of the cryptocurrency market. This truncated data window imposes a binding constraint on the empirical operationalization of the Block Maxima Method

(BMM). Given the limited sample size dictated by the younger stablecoins, this thesis faces a severe bias-variance trade-off. To retain a sufficient number of blocks for reliable maximum likelihood estimation of the GEV parameters the block size was restricted to a relatively narrow window of one month (30 days). A monthly block partition introduces notable statistical vulnerabilities. In financial asset returns, a one-month horizon risks violating the underlying assumptions of the BMM, as serial dependence or lingering volatility regimes can easily bridge across monthly boundaries. This means that the extracted Block Maxima may retain a degree of serial correlation, violating the independence assumption. Under unconstrained data conditions, such as multi-decade traditional fiat or equity series, it is standard practice to select an annual block size to guarantee asymptotic validity and independence [49, 21]. As an annual block size would reduce the stablecoin extreme value datasets to a handful of observations, making statistical inference impossible, the current framework represents a necessary but fragile compromise dictated entirely by the data limitations of the stablecoin asset class.

A related but distinct consequence of this lifespan asymmetry concerns the interpretability of the cross-pair dependence comparisons reported in Chapter 6. Because BTC and ETH are observed over a substantially longer and more cycle-complete history than USDT and USDC, the conclusion that the ETH–BTC pair exhibits significant tail dependence while pairs such as ETH–USDC do not is not a controlled comparison. The former dependence estimate is computed over the full 2015–2026 sample, spanning multiple bull and bear market regimes, whereas the latter is necessarily restricted to the shorter window over which USDC has existed, omitting several major cycles experienced by the volatile assets. Part of the contrast documented between same-volatility and cross-volatility pairs may therefore reflect differences in the length and market-regime composition of the underlying sample, rather than a purely structural difference in extremal dependence. This limitation does not change the qualitative conclusions so much, since the near-independence pattern between volatile assets and stablecoins is consistently observed across multiple pairs and under both the two-stage and full likelihood estimation approaches. However, it should moderate how strongly cross-pair dependence comparisons are interpreted, particularly those involving USDC.

The final limitation centers on the structural assumptions of the multivariate frameworks used to evaluate the safe haven hypothesis. In the previous chapter, a symmetric bivariate framework has been applied to a transformed dataset, reversing the sign of the volatile cryptocurrency returns to map severe market losses onto the positive right tail. These losses were then modeled jointly with the positive returns of the stablecoins using traditional Componentwise Maxima and threshold excess models. While this alignment strategy allows the application of standard bivariate EVT models, it introduces a conceptual mismatch with the underlying economic theory. Traditional bivariate models, such as the Logistic, Bilogistic or the Coles-Tawn, are fundamentally built to assess whether extreme events occur simultaneously in both components. However, the economic mechanism supporting a safe haven asset is conditional: the primary empirical concern is characterizing the behavior of the stablecoin given that the volatile asset has already entered a catastrophic tail state.

Instead of forcing this inherently directional question into a simultaneous joint tail framework, a more conceptually appropriate approach would be the deployment of advanced conditional extreme value models, most notably the framework established by Heffernan and Tawn [35]. The Heffernan-Tawn model bypasses the limitations of traditional multivariate EVT by directly capturing the conditional distribution of a response variable Y given that a conditioning variable X has exceeded an arbitrarily high threshold. Applied to this context, this methodology would allow for a direct, asymmetric evaluation of how stablecoins behave precisely during periods of extreme downside stress in the Bitcoin or Ethereum markets. By explicitly modeling the remaining dependence structure when one variable is extreme and the other is not, it would prevent the distortion of the dependence parameters ($\hat{\alpha}$ or $\hat{\chi}$) that occurs in simultaneous frameworks when joint extremes are naturally rare due to the bounded, flat nature of stablecoin pegging mechanisms.

Chapter 8

Conclusion

This thesis set out to investigate the tail risk of four cryptocurrencies, Bitcoin, Ethereum, Tether, and USD Coin, using the univariate and bivariate Extreme Value Theory framework defined by Coles [16], applied to two asset classes with markedly different volatility profiles: high-volatility cryptocurrencies (BTC, ETH) and stablecoins (USDT, USDC).

The first objective was to characterize the marginal tail behavior of each of the four assets using univariate EVT. The results have revealed a clear divide between volatile cryptocurrencies and stablecoins. In absolute terms, Bitcoin and Ethereum have exhibited substantially higher risk measures, return levels, VaR, and Expected Shortfall, than USDT and USDC. In terms of tail shape, however, the filtered residuals of the two stablecoins have provided evidence of Fréchet-type behavior, while those of Bitcoin and Ethereum have not offered sufficient grounds to reject a Gumbel-type specification. This apparent contrast reflects the structural nature of stablecoin de-pegging events, which represent disproportionately large shocks relative to their otherwise flat, low-volatility behavior, rather than evidence that stablecoins carry greater risk in absolute terms.

The second objective was to model the extremal dependence structure between all pairwise combinations of the four assets using bivariate EVT. Dependence between the two stablecoins, USDT and USDC, has proved weaker and not robust across model specifications, while pairs combining a volatile cryptocurrency with a stablecoin has showed no meaningful same-tail dependence under either the two-stage or full likelihood approach. Instead, a significant dependence is largely confined to the BTC–ETH pair. This dependence is consistently stronger in the upper tail than in the lower tail, meaning that the estimated parameter $\hat{\alpha}$ is smaller, for extreme gains than for extreme losses. Looking at these data, it is possible to conclude that for gains even more than for losses, a major portion of these extreme movements is not unique to Bitcoin but is equally true for Ethereum. This result provides a concrete justification for including Ethereum as a second high-volatility cryptocurrency, as it allows us to understand that in many cases, extreme gains are not just an isolated Bitcoin phenomenon, but rather a systemic market behavior that also involves Ethereum.

The third objective was to test explicitly for evidence of the safe haven hypothesis, by examining the cross-tail dependence between losses in volatile cryptocurrencies and gains in stablecoins. The asymmetric analysis of losses in BTC and ETH against gains in USDT and USDC has offered some evidence consistent with a flight to safety dynamic, although this evidence is limited in magnitude and needs to be interpreted as suggestive rather than conclusive. The hypothesis is therefore not rejected, but it cannot be considered robustly confirmed on the basis of the present analysis.

On the whole, these results indicate that EVT is a valuable framework to quantify tail risk in cryptocurrency markets and to distinguish the extremal behavior of volatile assets from that of stablecoins.

Future work could build on this analysis by addressing some of the limitations discussed in Chapter 7, especially those concerning the filtering stage used to obtain approximately i.i.d. residuals. The use of more advanced GARCH-type specifications may yield a more accurate filtering of the return series, potentially improving the performance of EVT in capturing the tail dynamics of these assets.

Bibliography

[1] Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of banking & finance*, 26(7):1487–1503, 2002.

[2] Rashad Ahmed, Iñaki Aldasoro, and Chanelle Duley. *Public information and stablecoin runs*. Bank for International Settlements, Monetary and Economic Department, 2024.

[3] Hirotugu Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 2003.

[4] Khamis Hamed Al-Yahyaee, Walid Mensi, and Seong-Min Yoon. Efficiency, multifractality, and the long-memory property of the bitcoin market: A comparative analysis with stock, currency, and gold markets. *Finance Research Letters*, 27:228–234, 2018.

[5] Andrew Ang and Allan Timmermann. Regime changes and financial markets. *Annu. Rev. Financ. Econ.*, 4(1):313–337, 2012.

[6] August A Balkema and Laurens De Haan. Residual life time at great age. *The Annals of probability*, pages 792–804, 1974.

[7] Dirk G Baur and Lai T Hoang. A crypto safe haven against bitcoin. *Finance Research Letters*, 38:101431, 2021.

[8] Dirk G Baur and Thomas K McDermott. Is gold a safe haven? international evidence. *Journal of Banking & Finance*, 34(8):1886–1898, 2010.

[9] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327, 1986.

[10] George Box and GM Jenkins. Analysis: Forecasting and control. *San francisco*, 10, 1976.

[11] Alessio Brini and Jimmie Lenz. A comparison of cryptocurrency volatility-benchmarking new and mature asset classes. *Financial Innovation*, 10(1):122, 2024.

[12] Pascal Bruhn and Dietmar Ernst. Assessing the risk characteristics of the cryptocurrency market: A garch-evt-copula approach. *Journal of Risk and Financial Management*, 15(8): 346, 2022.

[13] Kenneth P Burnham and David R Anderson. *Model selection and multimodel inference: a practical information-theoretic approach*. Springer, 2002.

[14] Vitalik Buterin et al. A next-generation smart contract and decentralized application platform. *white paper*, 3(37):2–1, 2014.

[15] Tom Celig, Tim Alvaro Ockenga, and Detlef Schoder. Distributional equality in ethereum? on-chain analysis of ether supply distribution and supply dynamics. *Humanities and Social Sciences Communications*, 12(1):1–15, 2025.

[16] Stuart Coles. *An introduction to statistical modeling of extreme values*, volume 208. Springer, 2001.

[17] Stuart G Coles and Jonathan A Tawn. Modelling extreme multivariate events. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 53(2):377–392, 1991.

[18] Laurens De Haan and Sidney I Resnick. Limit theory for multivariate sample extremes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 40(4):317–337, 1977.

[19] Lawrence T DeCarlo. On the meaning and use of kurtosis. *Psychological methods*, 2(3): 292, 1997.

[20] Paul Embrechts, Rüdiger Frey, Marius Hofert, and Alexander J. McNeil. Quantitative risk management – lecture slides. https://github.com/qrmtutorial/qrm/blob/master/slides/qrm.pdf.

[21] Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling extremal events: for insurance and finance*, volume 33. Springer Science & Business Media, 2013.

[22] Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007, 1982.

[23] Ronald Aylmer Fisher and Leonard Henry Caleb Tippett. Limiting forms of the frequency distribution of the largest or smallest member of a sample. In *Mathematical proceedings of the Cambridge philosophical society*, volume 24, pages 180–190. Cambridge University Press, 1928.

[24] Alexios Galanos. *rugarch: Univariate GARCH models.*, 2023. R package version 1.5-1.

[25] Ramazan Gençay, Faruk Selçuk, and Abdurrahman Ulugülyağci. High volatility, thick tails and extreme value theory in value-at-risk estimation. *Insurance: Mathematics and Economics*, 33(2):337–356, 2003.

[26] Manfred Gilli and Evis Këllezi. An application of extreme value theory for measuring financial risk. *Computational Economics*, 27(2):207–228, 2006.

[27] Konstantinos Gkillas and Paraskevi Katsiampa. An application of extreme value theory to cryptocurrencies. *Economics Letters*, 164:109–111, 2018.

[28] Konstantinos Gkillas and Costas Siriopoulos. Extreme correlation in cryptocurrency markets. *Available at SSRN 3180934*, 2018.

[29] Boris Gnedenko. Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of mathematics*, 44(3):423–453, 1943.

[30] Gary B Gorton and Jeffery Y Zhang. Taming wildcat stablecoins. *U. Chi. L. Rev.*, 90:909, 2023.

[31] Gordon Gudendorf and Johan Segers. Extreme-value copulas. In *Copula Theory and Its Applications: Proceedings of the Workshop Held in Warsaw, 25-26 September 2009*, pages 127–145. Springer, 2010.

[32] Emil J Gumbel. Bivariate logistic distributions. *Journal of the American Statistical Association*, 56(294):335–349, 1961.

[33] Peter R Hansen and Asger Lunde. A forecast comparison of volatility models: does anything beat a garch (1, 1)? *Journal of applied econometrics*, 20(7):873–889, 2005.

[34] Fredrik Hedman and Emil Håkansson. Estimating expected shortfall using parametric and non-parametric approaches. 2020.

[35] Janet E Heffernan and Jonathan A Tawn. A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(3):497–546, 2004.

[36] Eric Hillebrand. Neglecting parameter changes in garch models. *Journal of Econometrics*, 129(1-2):121–138, 2005.

[37] Marius Hofert, Rüdiger Frey, and Alexander J McNeil. The quantitative risk management exercise book, 2020.

[38] Jonathan Richard Morley Hosking, James R Wallis, and Eric F Wood. Estimation of the generalized extreme-value distribution by the method of probability-weighted moments. *Technometrics*, 27(3):251–261, 1985.

[39] Jonathan RM Hosking. Algorithm as 215: Maximum-likelihood estimation of the parameters of the generalized extreme-value distribution. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 34(3):301–310, 1985.

[40] Cho-Hoi Hui, Andrew Wong, and Chi-Fai Lo. Stablecoin price dynamics under a peg-stabilising mechanism. *Journal of International Money and Finance*, 152:103280, 2025.

[41] Harry Joe, Richard L Smith, and Ishay Weissman. Bivariate threshold methods for extremes. *Journal of the royal statistical society: series B (methodological)*, 54(1):171–183, 1992.

[42] Paraskevi Katsiampa. Volatility estimation for bitcoin: A comparison of garch models. *Economics letters*, 158:3–6, 2017.

[43] Hanbeen Kim, Sooyoung Kim, Hongjoon Shin, and Jun-Haeng Heo. Appropriate model selection methods for nonstationary generalized extreme value models. *Journal of Hydrology*, 547:557–574, 2017.

[44] Dimitrios Koutmos. Return and volatility spillovers among cryptocurrencies. *Economics Letters*, 173:122–127, 2018.

[45] Anthony W Ledford and Jonathan A Tawn. Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1):169–187, 1996.

[46] Aleksandr Michajlovič Ljapunov. *Nouvelle forme du théorème sur la limite de probabilité*. 1901.

[47] Greta M Ljung and George EP Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978.

[48] Francois Longin and Bruno Solnik. Extreme correlation of international equity markets. *The journal of finance*, 56(2):649–676, 2001.

[49] François M Longin. The asymptotic distribution of extreme stock market returns. *Journal of business*, pages 383–408, 1996.

[50] Richard K Lyons and Ganesh Viswanath-Natraj. What keeps stablecoins stable? *Journal of International Money and Finance*, 131:102777, 2023.

[51] Alexander J McNeil and Rüdiger Frey. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of empirical finance*, 7(3-4):271–300, 2000.

[52] Alexander J McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.

[53] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.

[54] Roger B Nelsen. *An introduction to copulas*. Springer, 2006.

[55] Pitipat Nittayakamolphun, Thanchanok Bejrananda, and Panjamapon Pholkerd. Stablecoins as safe haven or hedging asset for cryptocurrencies (in thai). *Asian Journal of Applied Economics/Applied Economics Journal*, 29(2):45–70, 2022.

[56] Myung Hyun Park and Joseph HT Kim. Estimating extreme tail risk measures with generalized pareto distribution. *Computational Statistics & Data Analysis*, 98:91–104, 2016.

[57] Yudi Pawitan. *In all likelihood: statistical modelling and inference using likelihood*. Oxford University Press, 2001.

[58] James Pickands III. Statistical inference using extreme order statistics. *the Annals of Statistics*, pages 119–131, 1975.

[59] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2022. URL https://www.R-project.org/.

[60] R Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of banking & finance*, 26(7):1443–1471, 2002.

[61] Abhay K Singh, David E Allen, and Powell J Robert. Extreme market risk and extreme value theory. *Mathematics and computers in simulation*, 94:310–328, 2013.

[62] M Sklar. Fonctions de répartition à n dimensions et leurs marges. In *Annales de l'ISUP*, volume 8, pages 229–231, 1959.

[63] A. G. Stephenson. evd: Extreme value distributions. *R News*, 2(2):31–32, June 2002. URL https://CRAN.R-project.org/doc/Rnews/.

[64] Ruey S Tsay. *Analysis of financial time series*. John wiley & sons, 2005.

[65] Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

[66] Dimitri Vayanos. Flight to quality, flight to liquidity, and the pricing of risk, 2004.

[67] Abraham Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical society*, 54 (3):426–482, 1943.

[68] Abraham Wald and Jacob Wolfowitz. On a test whether two samples are from the same population. *The Annals of Mathematical Statistics*, 11(2):147–162, 1940.

[69] Evert Wipplinger. Philippe jorion: value at risk-the new benchmark for managing financial risk. *Financial Markets and Portfolio Management*, 21(3):397, 2007.

[70] Yasuhiro Yamai and Toshinao Yoshiba. Value-at-risk versus expected shortfall: A practical perspective. *Journal of banking & finance*, 29(4):997–1015, 2005.

Appendix A

GEV Diagnostic Plots

This appendix collects the remaining GEV diagnostic plots not displayed in Section 5.1.3. For USDC and USDT, the plots for losses are presented, showing a qualitatively similar but slightly less severe pattern of poor tail fit compared to the figures in the main text. The diagnostic plots for Bitcoin and Ethereum are also included, where the Block Maxima GEV model performs considerably better, with no systematic deviations in the tail region.

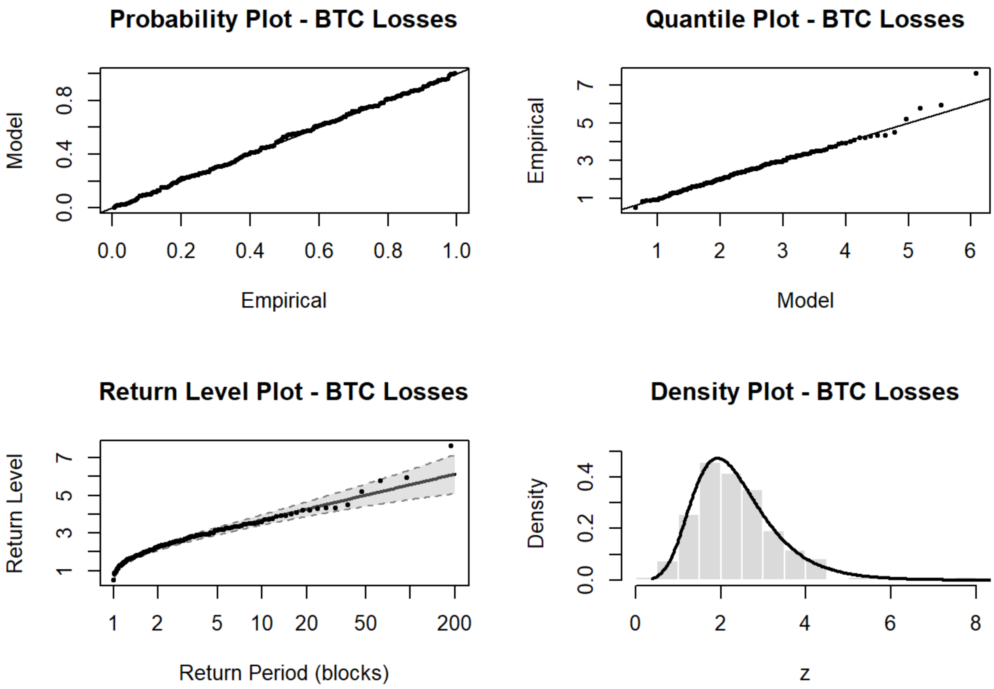

Figure A.1: GEV diagnostic plots for BTC losses. The probability plot shows observations closely aligned with the diagonal, and the quantile and return level plots indicate strong agreement between empirical and theoretical quantiles, with the majority of Block Maxima falling within the confidence bands.

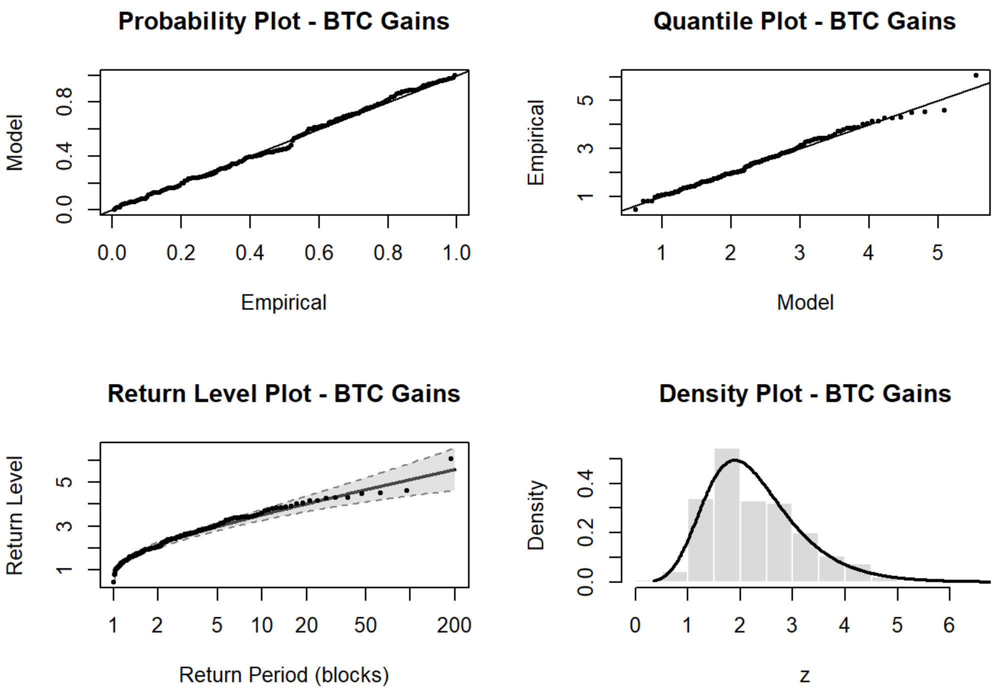

Figure A.2: GEV diagnostic plots for BTC gains. The fit is marginally weaker than for losses, with minor deviations in the upper tail of the quantile plot, but the density plot confirms that the fitted distribution closely matches the empirical histogram and the overall fit remains adequate.

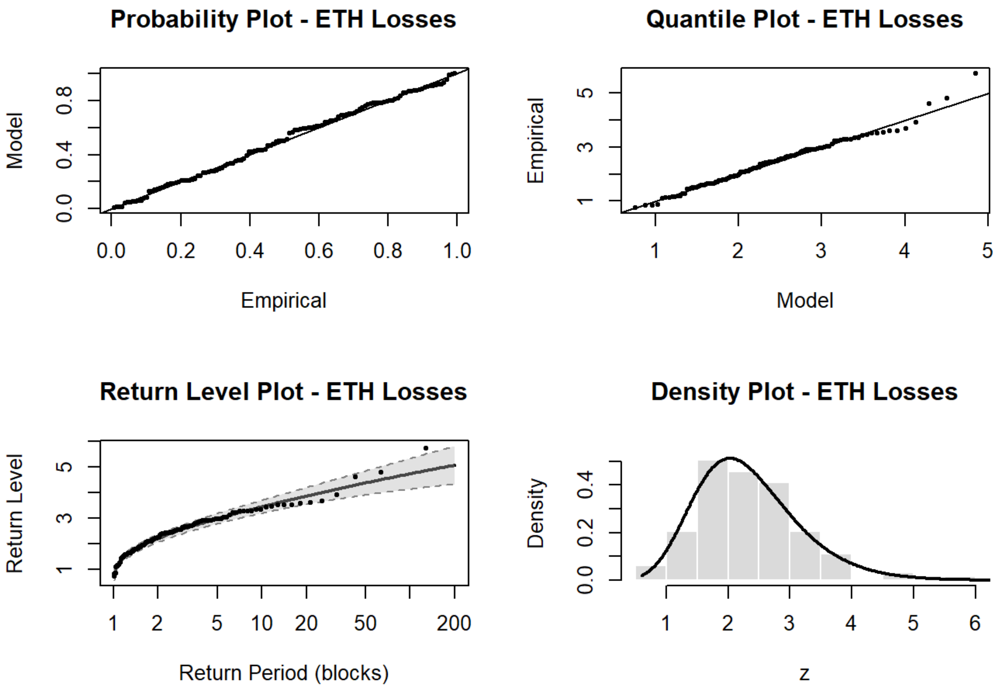

Figure A.3: GEV diagnostic plots for ETH losses. The probability and quantile plots display no clear evidence of systematic misspecification, and the model captures tail behavior reasonably well, with discrepancies limited in magnitude.

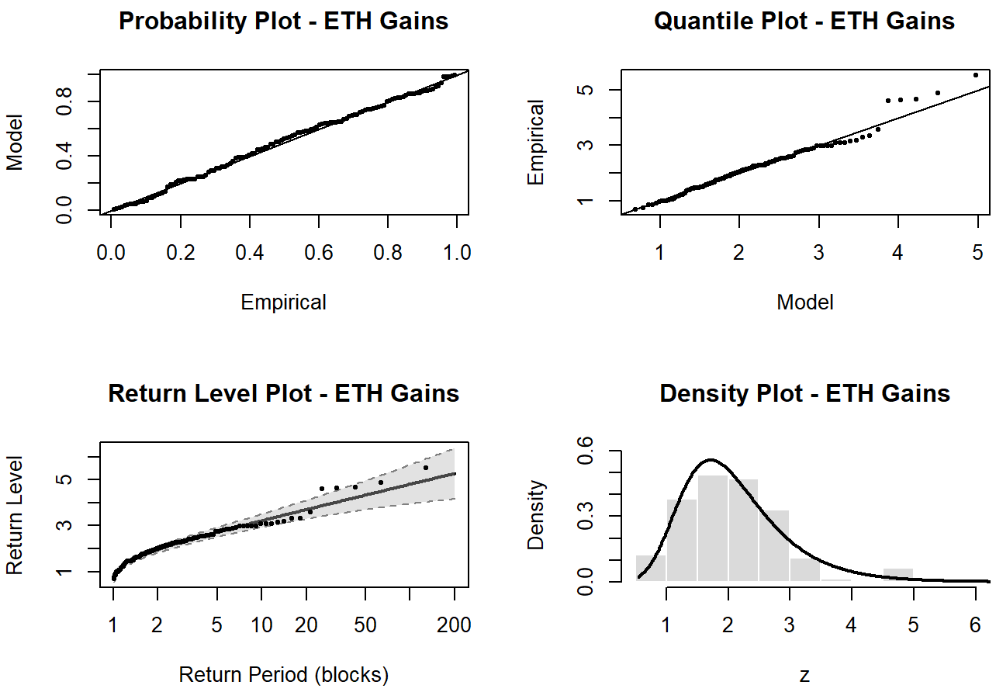

Figure A.4: GEV diagnostic plots for ETH gains. Some deviation is visible in the upper tail of the quantile plot, and a small number of observations fall outside the return level confidence bands, but these discrepancies do not materially affect the overall adequacy of the model.

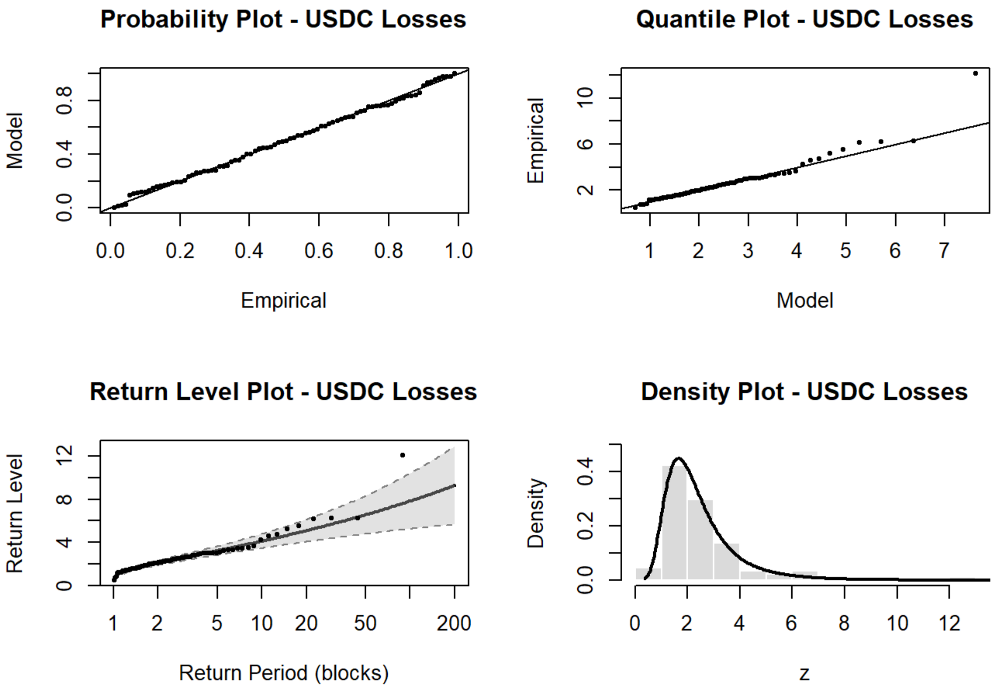

Figure A.5: GEV diagnostic plots for USDC losses. The density and probability plots show a cleaner fit than the gains case (Figure 5.5), but the quantile and return level plots reveal that one outlier still slips outside the 95% confidence band, reflecting the same structural issue in a less pronounced form.

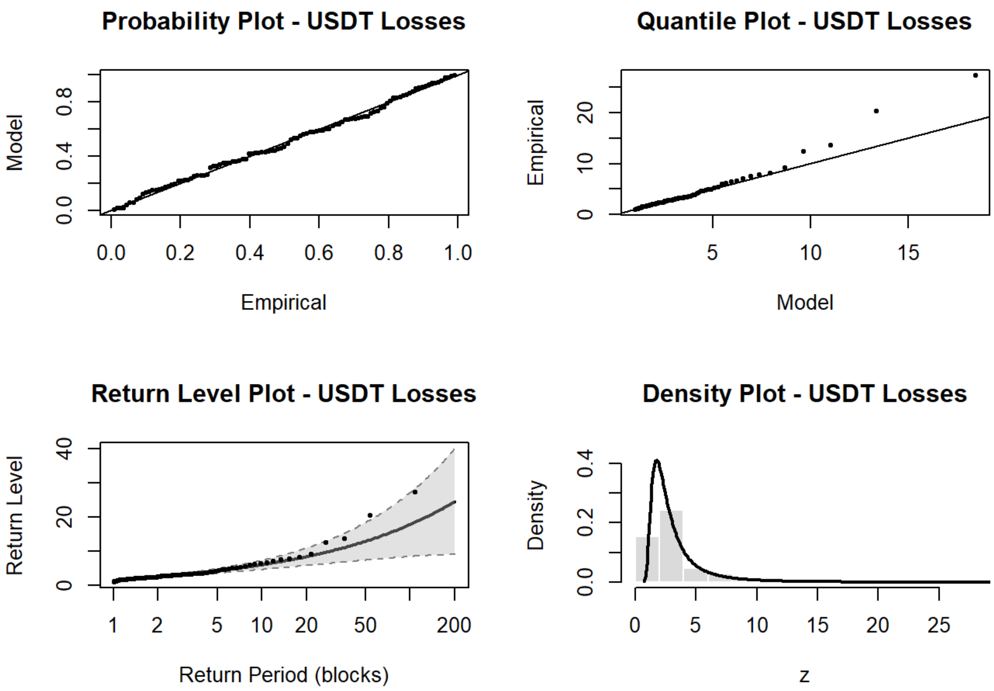

Figure A.6: GEV diagnostic plots for USDT losses. Deviations are visible in the probability and density plots, and the quantile and return level plots display suboptimal results. Unlike the gains case (Figure 5.6), no single outlier dramatically breaches the 95% confidence band, yet the overall fit remains far from optimal.

Appendix B

POT Diagnostic Plots

This appendix collects the remaining POT diagnostic plots not displayed in Section 5.2.4. For USDC and USDT, the gain-side and loss-side plots, respectively, reveal a qualitatively similar pattern of poor tail fit to that documented in the main text, though slightly less severe. The diagnostic plots for Bitcoin and Ethereum, by contrast, suggest that the POT–GPD model performs considerably better (both relative to the stablecoin pairs and relative to the GEV-based analysis of the same assets), with no systematic departures from the fitted distribution in the tails.

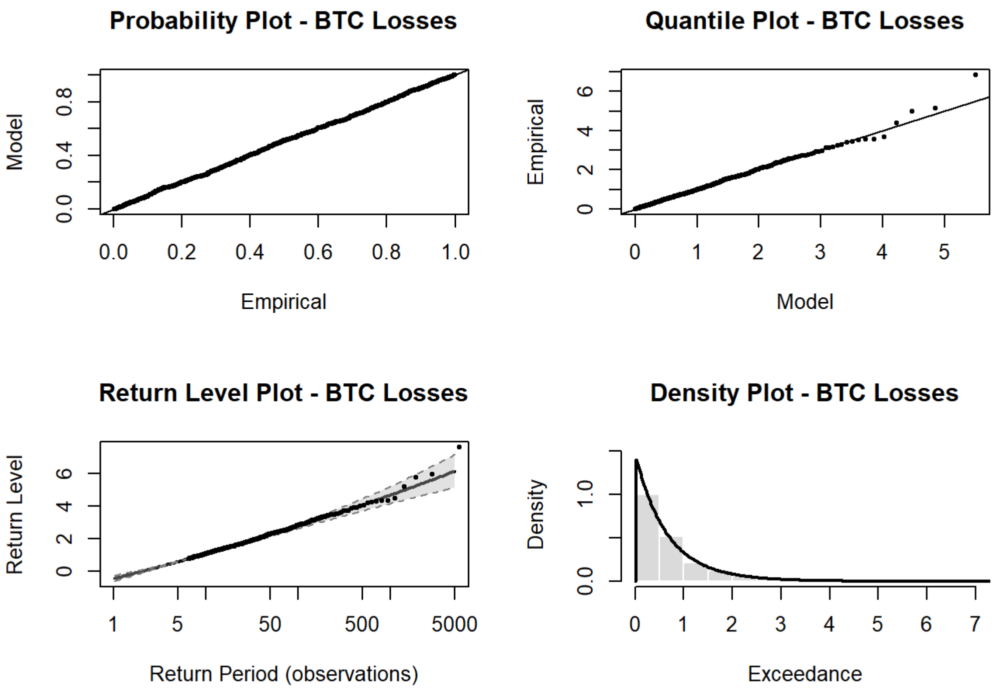

Figure B.1: GPD diagnostic plots for BTC losses. The probability plot shows observations closely aligned with the diagonal, and the quantile and return level plots indicate a good fit, with only minor deviations in the very upper tail.

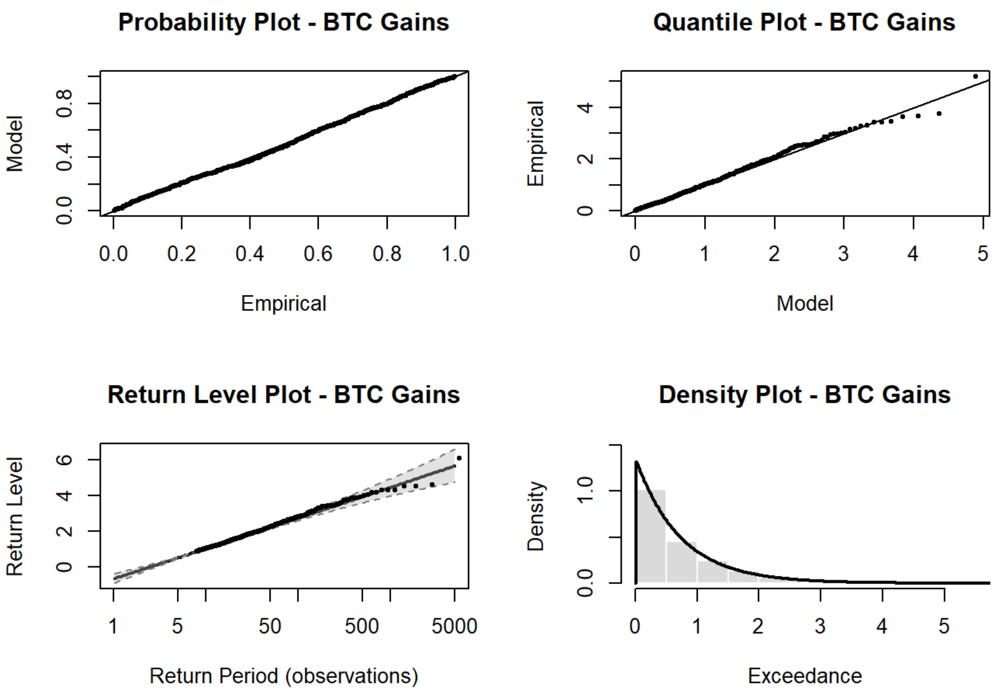

Figure B.2: GPD diagnostic plots for BTC gains. The fit is marginally worse than for losses, with somewhat more pronounced upper tail deviations, but the model provides an adequate description of the tail behavior.

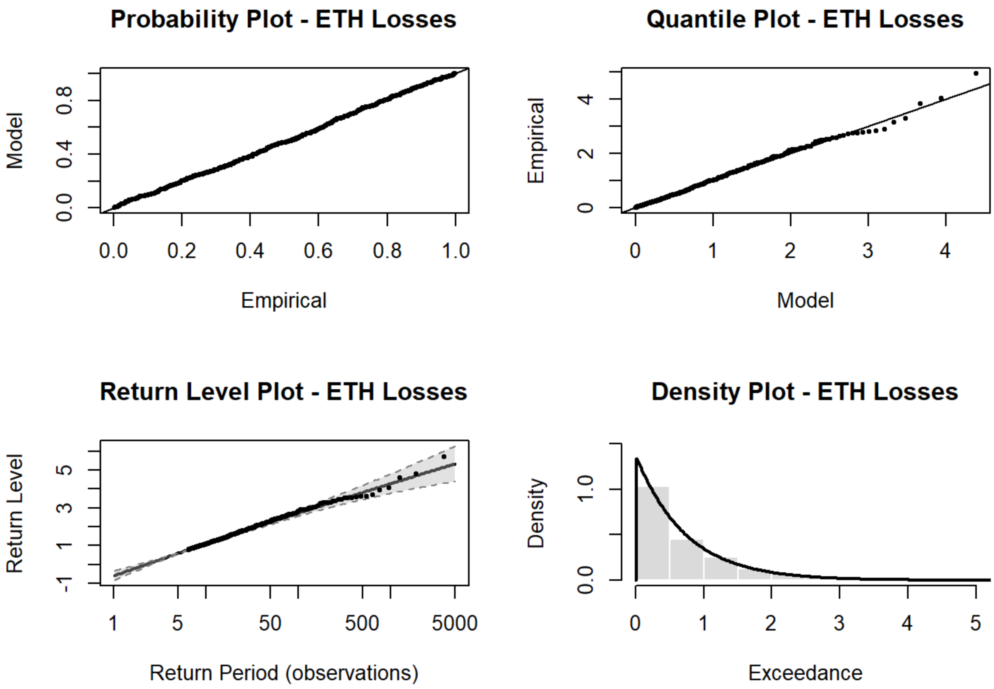

Figure B.3: GPD diagnostic plots for ETH losses. The probability and quantile plots display no clear evidence of systematic misspecification, and the model captures tail behavior reasonably well, with discrepancies limited in magnitude.

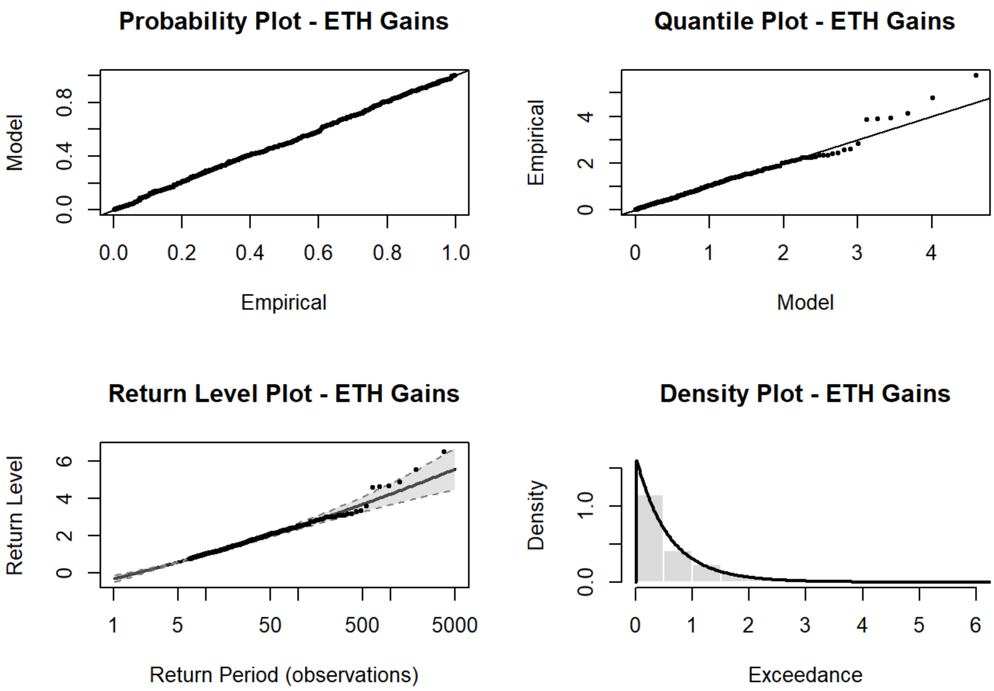

Figure B.4: GPD diagnostic plots for ETH gains. Some deviation is visible in the upper tail of the quantile plot and a small number of exceedances fall outside the return level confidence bands, but these discrepancies do not materially affect the overall adequacy of the model.

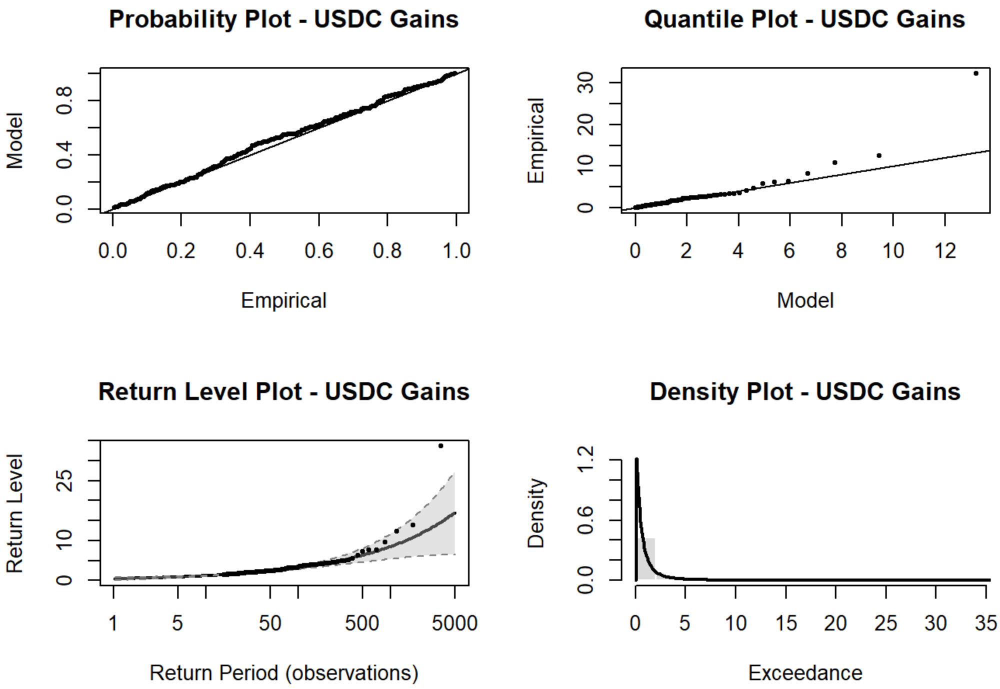

Figure B.5: GPD diagnostic plots for USDC gains. A clear improvement over the GEV fit (Figure 5.5), with substantially better alignment in the probability and quantile plots, though some difficulty in the extreme upper tail remains, such as in the losses case (Figure 5.14).

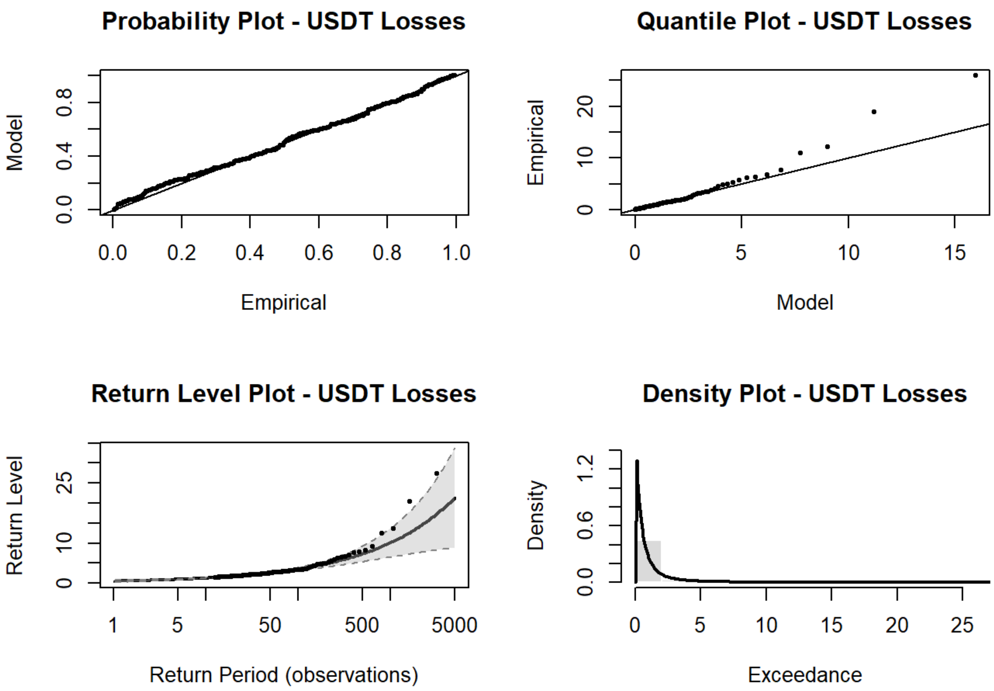

Figure B.6: GPD diagnostic plots for USDT losses. The probability plot shows a marginal improvement in the central fit relative to the GEV model, but the quantile and return level plots continue to exhibit large outliers breaching the confidence bands, mirroring the pattern observed for USDT gains (Figure 5.13).

Appendix C

Visual Independence-Bivariate Threshold Approach

This appendix collects the remaining χ and $\bar{\chi}$ plots from Section 6.2.2. Only the mixed pairs (BTC/ETH with a stablecoin) are reported here, as the homogeneous pairs (stablecoin–stablecoin and high-volatility–high-volatility) are already discussed in the main chapter. As suggested by the $\hat{\alpha}$ estimates and confirmed by the plots below, these mixed pairs are nearly independent: both χ and $\bar{\chi}$ consistently indicate the absence of asymptotic dependence, for gains and losses alike. Throughout, blue distinguishes the high-volatility assets (BTC/ETH) from the stablecoins, consistent with the color convention used in the corresponding main chapter.

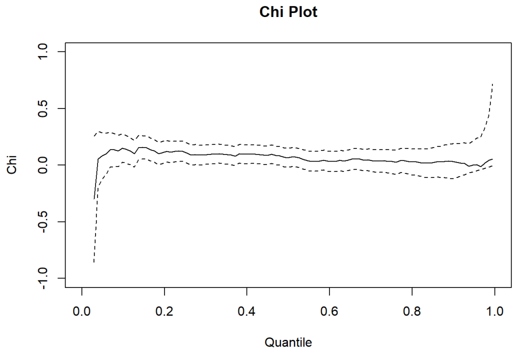

(a) BTC–USDT χ (gains)

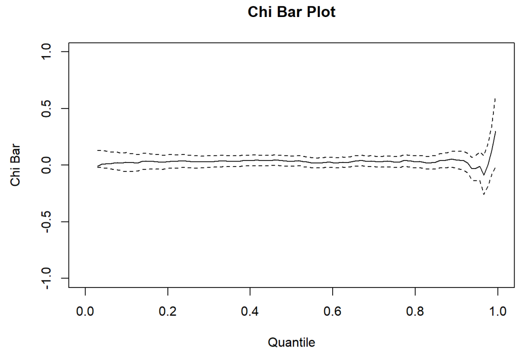

(b) BTC–USDT $\bar{\chi}$ (gains)

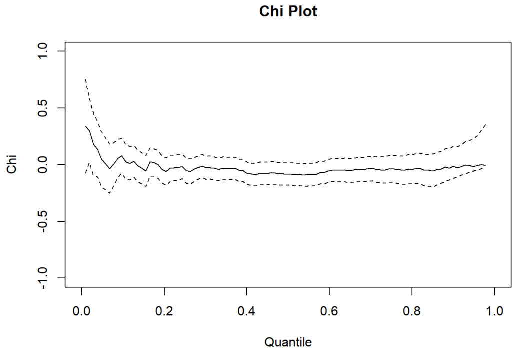

(c) BTC–USDC χ (gains)

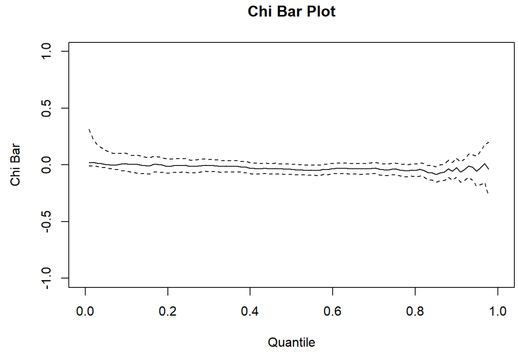

(d) BTC–USDC $\bar{\chi}$ (gains)

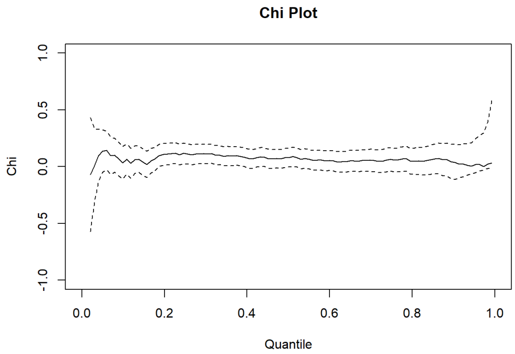

(e) ETH–USDT χ (gains)

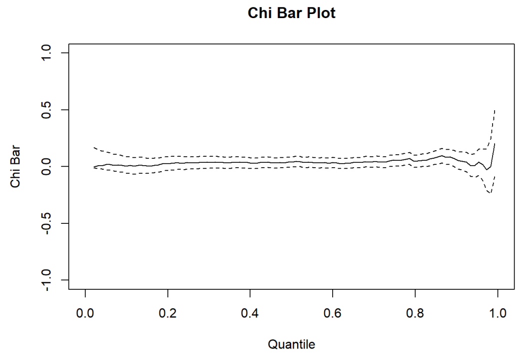

(f) ETH–USDT $\bar{\chi}$ (gains)

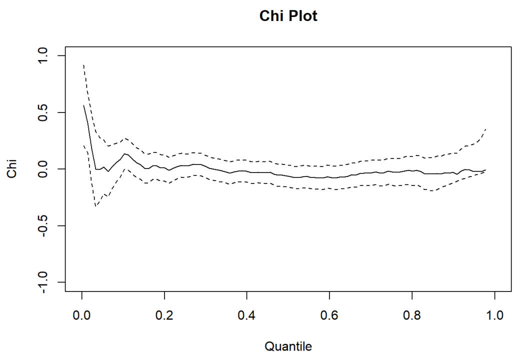

(g) ETH–USDC χ (gains)

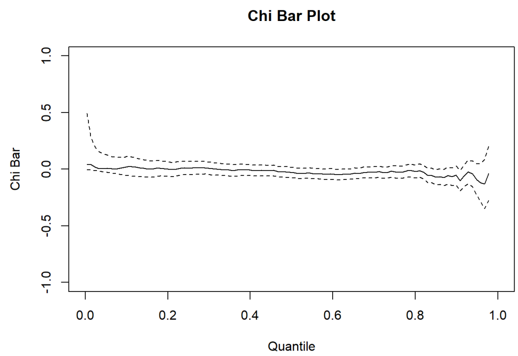

(h) ETH–USDC $\bar{\chi}$ (gains)

Figure C.1: χ and $\bar{\chi}$ plots for the remaining pairs (gains). No asymptotic dependence is detected.

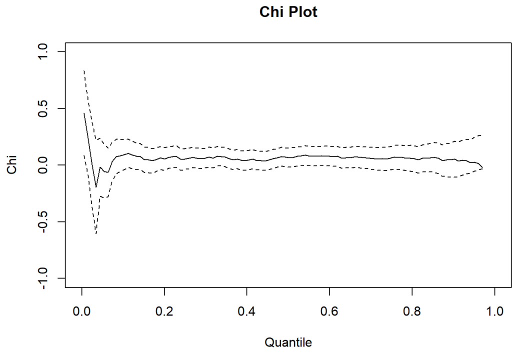

(a) BTC–USDT χ (losses)

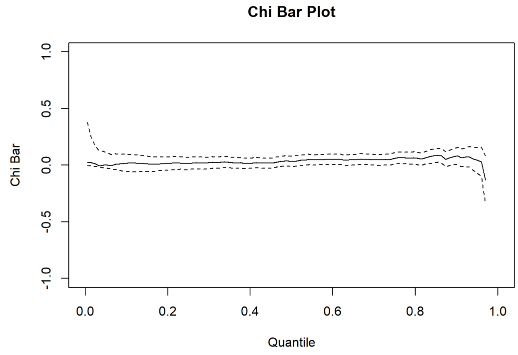

(b) BTC–USDT $\bar{\chi}$ (losses)

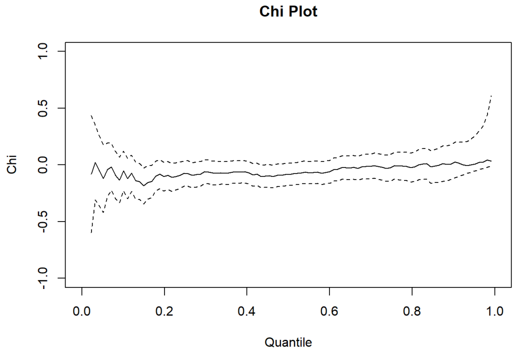

(c) BTC–USDC χ (losses)

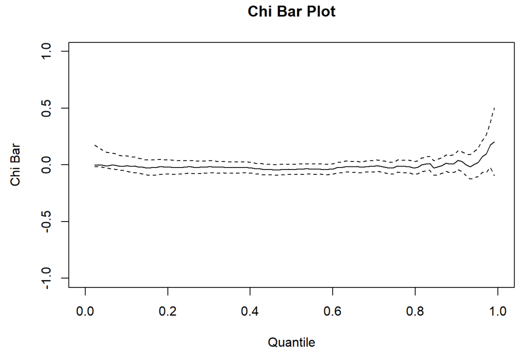

(d) BTC–USDC $\bar{\chi}$ (losses)

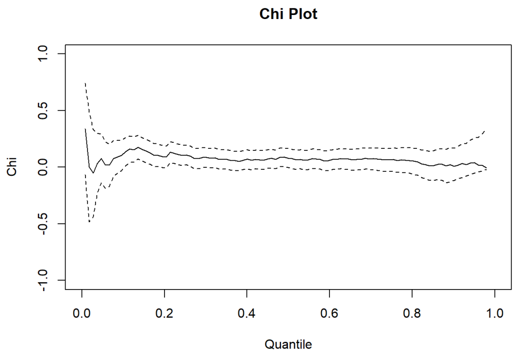

(e) ETH–USDT χ (losses)

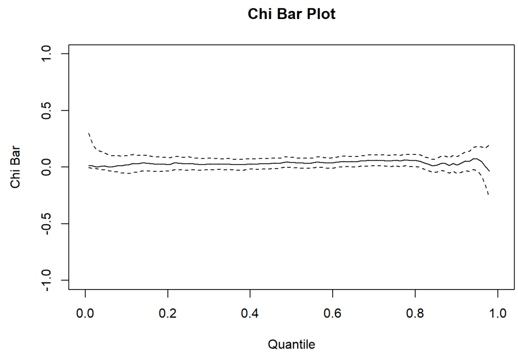

(f) ETH–USDT $\bar{\chi}$ (losses)

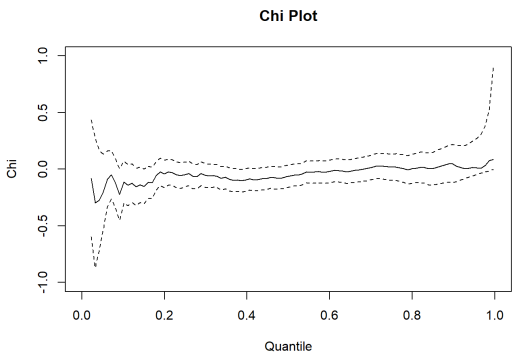

(g) ETH–USDC χ (losses)

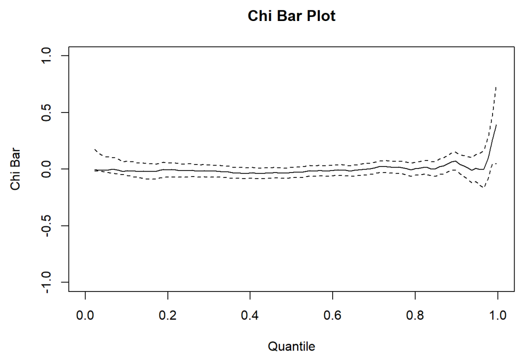

(h) ETH–USDC $\bar{\chi}$ (losses)

Figure C.2: χ and $\bar{\chi}$ plots for the remaining pairs (losses). No asymptotic dependence is detected.

Acknowledgements

I wish to express my sincere gratitude to my supervisor, Professor Ilaria Prosdocimi, for her invaluable guidance and support.

To my mother, Ingrid, who has, in her own way, taught me the meaning of perseverance.

To my sister, Sofia, who reminds me every day that growth begins with questioning oneself, but never with doubting one's worth.

This achievement is due, in no small part, to their support.

Grazie,
Lisa