



Tail Risk Modelling in Foreign Exchange Markets: Evidence from Botswana Pula, Bitcoin, and South African Rand

Edwin Moyo^{1,3*}, Katleho Makatjane², Ramasamy Sivasamy¹

¹*Department of Statistics, University of Botswana, Botswana*

²*Department of Statistics and Population Studies, University of the Western Cape, South Africa*

³*Department of Mathematics and Statistics, Mulungushi University, Kabwe, Zambia*

Abstract The study uses data for each currency pair (Bitcoin vs USD, South African Rand vs USD, and Botswana Pula vs USD) to model the extreme quantiles of their weekly returns using two distributional frameworks: the Generalised Extreme Value (GEV) and the Generalised Pareto Distribution (GPD). Two methodological approaches are employed to estimate these distributions: the block maxima/minima approach and the peaks-over-threshold (POT) approach. To ensure optimal estimation of the parameters of the distributions, we employ a stability-prediction threshold optimisation algorithm. The predictive error metrics (mean absolute error, mean squared error, root mean squared error) and maximum likelihood estimation (MLE) are utilised to provide unbiased estimates of the distribution parameters. In addition, goodness-of-fit diagnostics and the information criteria (Akaike Information Criterion and Bayesian Information Criterion) are employed to evaluate the appropriateness of each distributional framework for modelling the tails of the three currencies' distributions. Our empirical evidence shows important differences among the three currencies in their tail behaviour. Specifically, the findings show that Bitcoin/USD has pronounced heavy-tailed behaviour relative to both ZAR/USD and BWP/USD. Furthermore, the GEV was found to accurately model the extreme gains experienced by investors holding Bitcoin, whereas the GPD better captures extreme losses on this asset. On the other hand, the GPD is found to be a much better model than the GEV for capturing moderate tail risk in ZAR/USD, especially for extreme depreciations against the US dollar. Finally, the BWP/USD exhibits very little heavy-tailed behaviour, that is, thin tails, likely due to its more stable currency regime compared to either ZAR/USD or BTC/USD. Overall, empirical results show that the GPD is a much more reliable predictor of both upside and downside extreme events for all three currency pairs examined in this paper.

Keywords Efficiency, Exchange rate return, Extreme Value Theorem, Generalised Pareto Distribution, Generalised Extreme Distribution, Block Maxima, Peak over Threshold

AMS 2010 subject classifications 62-07, 62M10

DOI: 10.19139/soic-2310-5070-????

1. Introduction

Capturing extreme and sudden events, namely tail risks, in financial markets is essential for understanding and mitigating the impact of exogenous shocks. Foreign exchange markets, in particular, present significant modelling challenges as they are frequently subject to large and unanticipated fluctuations. Empirical evidence from [1] and [2] suggests that exchange rate movements are often accompanied by abrupt shifts in volatility, making reliable forecasting and risk assessment particularly arduous. Consequently, accurate modelling requires a precise specification of the statistical properties of return distributions, as distributional assumptions directly dictate the validity of inference, forecasting performance, and the quantification of financial risk.

Linear models, such as the Autoregressive Integrated Moving Average (ARIMA) and its variants, along with exponential smoothing techniques, have long been used in financial analysis; however, these are fundamentally

*Correspondence to: Edwin Moyo (Email: emoyo@mu.edu.zm, Department of Statistics, University of Botswana, Botswana.

constrained by their inability to capture the complex, non-linear dynamics of modern assets like Bitcoin. As noted by [3], the linear architecture of ARIMA fails to account for the asymmetries and structural breaks inherent in volatile markets. This limitation necessitated a shift toward models capable of addressing non-constant variance; hence, the Autoregressive Conditional Heteroscedasticity (ARCH) model, introduced by [4] which allowed the variance to depend on past squared residuals. Yet ARCH models often require a large number of parameters to capture volatility persistence, thereby compromising estimation efficiency. The Generalised Autoregressive Conditional Heteroscedasticity (GARCH) model by [5] addressed this by incorporating lagged conditional variance in a more parsimonious manner. Nevertheless, standard GARCH specifications assume symmetric volatility responses and thus fail to capture the leverage effect, in which negative shocks exert a more pronounced influence on volatility than positive shocks of similar magnitude. This gave rise to asymmetric extensions, including the Exponential GARCH (EGARCH) model by [6], which utilises logarithmic variance to ensure positivity and capture asymmetry, and the GJR-GARCH model of [7]. Similarly, the Threshold GARCH (TGARCH) model of [8] shifted the focus to conditional standard deviation to better reflect the magnitude of price shocks.

Despite these developments, the inherent asymmetry and heavy-tailed nature of financial time series data often limit GARCH-type models' ability to effectively capture the underlying time-varying process. To mitigate this, [9] enhanced these architectures by introducing regime-switching models, known as Markov-switching GARCH (MS-GARCH). These models permit the simultaneous modelling of distinct high and low volatility states while capturing regime-specific asymmetry. However, a critical methodological bottleneck remains. The MS-GARCH models are frequently constrained by assumptions of normality or standard parametric shapes, which are insufficient to capture the tails of the distribution. While regime-switching can account for structural shifts, it often fails to accurately quantify the magnitude of rare, catastrophic events that lie beyond the model's expected variance within a given state.

These persistent limitations emphasise the need for approaches that focus explicitly on extreme observations, leading to the adoption of Extreme Value Theory (EVT). As identified in recent studies [10, 11], conditional heteroscedasticity models, including MS-GARCH, often struggle to capture the extreme leptokurtosis and long-memory properties observed in volatile assets. Non-parametric approaches, such as historical simulation (HS), face similar hurdles. While HS avoids strict distributional assumptions, its reliance on the independent and identically distributed (i.i.d.) assumption makes it prone to ghost effects and slow adaptation to sudden market shifts, as emphasised by [12, 13]. To overcome this, recent studies have shifted to the EVT methodology, which has since become a prominent area for modelling rare but impactful events. Specifically, the Generalised Extreme Value (GEV) distribution, suitable for modelling block maxima or minima [14], and the Generalised Pareto Distribution (GPD), applied within the Peaks Over Threshold (POT) approach, provide a robust basis for capturing extreme movements. However, the practical implementation of EVT presents a significant methodological dilemma; the selection of the block size for GEV and the threshold for GPD. Currently, threshold selection often relies on graphical tools, such as mean excess plots, which [15] and [16] argue are inherently subjective and difficult to replicate. Similarly, block size selection often follows arbitrary calendar conventions rather than statistical suitability. As noted by [17], inappropriate block sizes lead to a critical trade-off where small blocks introduce bias from non-extreme values, while excessively large blocks increase estimator variance.

With this said, we address these problems in this study by developing two complementary, data-driven procedures within the EVT structure. First, we propose a systematic algorithm to determine the optimal block size for the GEV distribution, utilising independence diagnostics, goodness-of-fit tests, and predictive metrics such as mean square error (MSE), mean absolute error (MAE), and root mean square error (RMSE). Second, we develop a Stability-Prediction Threshold Optimisation Algorithm (SPTOA) for the POT approach. By minimising a composite objective function that balances parameter stability with predictive accuracy, the SPTOA replaces subjective graphical interpretation with a formal optimisation criterion. The robustness of the proposed procedures lies in their transition from heuristic-based estimation to formal optimisation. By utilising a systematic algorithm for GEV block-size selection, we effectively resolve the bias-variance trade-off that often plagues extreme value modelling. This data-driven approach ensures that the extracted maxima or minima represent true extreme events while maintaining sufficient sample size for reliable inference. Furthermore, the SPTOA introduces a mathematically rigorous alternative to the subjective graphical methods commonly used in the POT approach.

By minimising a composite objective function that simultaneously prioritises parameter stability and predictive accuracy, the algorithm provides a transparent, replicable, and robust mechanism for tail-risk estimation in highly volatile foreign exchange markets.

To formalise and validate these algorithms, this research examines the tail behaviour of three distinct exchange rate series: Bitcoin/USD (BTC/USD), the South African Rand/USD (ZAR/USD), and the Botswana Pula/USD (BWP/USD). These assets represent a spectrum of market volatility: Bitcoin as a speculative digital asset, the Rand as a liquid but volatile emerging-market currency, and the Pula as a relatively stable regional currency pegged to a basket inclusive of the Rand [18]. The contribution of this paper to the financial modelling and extreme value theory literature is fourfold. First, we provide a data-driven algorithm for optimal block-size selection, filling a long-standing methodological void in block maxima analysis by replacing arbitrary calendar conventions with statistical optimisation. Second, we introduce the SPTOA to formalise threshold selection, replacing subjective graphical heuristics with a formal optimisation criterion that enhances the transparency and reproducibility of GPD modelling. Third, a rigorous comparative analysis of GEV and GPD distributions across a speculative digital asset BTC/USD, a liquid emerging-market currency ZAR/USD, and a stable regional currency (BWP/USD) is provided, offering new insights into the disparate tail-risk dynamics of these financial systems. Finally, a joint backtesting framework for Value-at-Risk and Expected Shortfall is employed to evaluate the predictive integrity of the proposed models. By utilising joint loss functions and encompassing tests, we ensure that the tail-risk estimates for all three exchange rate pairs are not only statistically robust but also practically reliable for institutional risk management.

Research Highlights and Key Findings

This study advances the literature by applying Extreme Value Theory (EVT) models, namely the Generalised Extreme Value (GEV) distribution and the Generalised Pareto Distribution (GPD), to capture the extreme quantiles of exchange rate risk in the BWP/USD, ZAR/USD, and BTC/USD markets. The analysis distinguishes between downside risk (losses) and upside risk (gains), thereby providing a comprehensive view of tail behaviour in both directions. By comparing model fit, shape parameters, and goodness-of-fit tests, this research identifies which model more effectively characterises the extremal dynamics of each currency, offering critical insights for traders, policymakers, and risk managers. The highlights and key findings are summarised in Table 1.

Table 1. Performance and Contributions of the GEV and GPD Models in Exchange Rate Risk

Characteristic	Highlights	Findings	Contributions
Model Fit – BTC/USD	GEV provides an excellent fit for maxima	Significant and positive EVI indicates heavy right tails	Fréchet distribution is suitable for extreme gains
	GPD fits minima reliably	EVI for losses is not significant, implying exponential-type tails	GPD (exponential distribution) captures downside risk more effectively
Model Fit – ZAR/USD	GEV indicates heavy tails but rejected by goodness-of-fit tests	Positive and significant EVI values for both gains and losses	Highlights the need for GPD in modelling extremes
	GPD fits losses with bounded tails	Negative and significant EVI for downside risk	Pareto type II distribution captures extreme losses
Model Fit – BWP/USD	Both GEV and GPD yield non-significant EVIs	No strong evidence of heavy tails in either gains or losses	Gumbel and exponential distributions best describe extremes
	GPD demonstrates strong statistical fit	All goodness-of-fit tests (KS, CVM, AD) validate performance	GPD superior for both maxima and minima
Comparative Risk	BTC/USD exhibits heavier tails in maxima compared to BWP/USD and ZAR/USD	Higher ξ values imply greater risk exposure for Bitcoin	Identifies BTC/USD as the riskiest in terms of upside extremes, followed by ZAR/USD, while BWP/USD is least risky
Policy Implications	EVT captures tail risks overlooked by volatility models	Central banks can enhance stability by incorporating EVT into monitoring	provide a framework for better preparedness against exchange rate shocks

This study is organised as follows: section 2 presents the methodology. In Section 4, a preliminary analysis of the data is conducted to investigate its nature, and Section 5 is devoted to the conclusion.

2. Methodology

This section outlines the data sources, modelling framework, estimation procedures, and validation techniques employed in this study. The daily currency data from January 2, 2015, to January 31, 2026, is obtained from [Yahoo finance website](#) and [Bank of Botswana website](#). The statistical analysis is conducted in R [19] using RStudio [20]. Extreme value modelling is implemented using the `ismev` package of [21], the `evd` package of [22], the `evir` package of [23], and the `extRemes` package of [24].

Bitcoin is traded continuously in cryptocurrency markets, resulting in 3083 daily observations. In contrast, the South African rand and Botswana pula are not traded on weekends and public holidays, resulting in 2621 observations for each series. For effective analysis, [16] argue that missing values for ZAR/USD and BWP/USD should be replaced with zero (0) returns, which indicate neither a gain nor a loss during periods when the markets are closed. This treatment ensures a consistent time index across the return series and avoids distortions that may arise from interpolation or removal of non-trading days. Consequently, the study is conducted using a complete return series.

2.1. Extreme Value Theory Approach to Modelling Downside and Upside Risk

This study uses Extreme Value Theory, which is based on previous studies showing its effectiveness for modelling the tails of distributions, where extreme values are important to model. EVT examines the probabilistic behaviour of extreme values in a distribution, as opposed to modelling the centre of the distribution. There are two common methods used in EVT: the Block Maxima Method and the Peaks Over Threshold (POT) method. Both methods were used in this study, as they have been used in other studies of extreme exchange rate fluctuations, such as [26]. The Block Maxima Method models the maximum (or minimum) value of each "block" of data. The POT method models the number of times an observation exceeds a certain threshold (above a sufficiently high threshold) using the Generalised Pareto Distribution (GPD). The Block Maxima method is theoretically justified by the Fisher-Tippett-Gnedenko theorem [27], which states that the limiting distribution of properly normalised maxima will converge to the Generalised Extreme Value (GEV) distribution. The POT method is also theoretically justified by the Pickands-Balkema-de Haan theorem [28], which states that exceedances above a sufficiently high threshold will converge to the GPD. The conclusions of these two theorems provide the theoretical basis for modelling extremes using the GEV and GPD distributions.

2.2. Generalised Extreme Value Distribution and BMM

The block maxima/minima method is a popular approach to EVT. The block maxima/minima method (BMM) identifies extreme values by extracting the maximum and minimum values from non-overlapping blocks of data. As [29] indicates, the GEV distribution is typically fitted to these maxima/minima. The BMM method [25] identifies extreme values within a block of data, and one major drawback of using the BMM is the need to select a suitable block size. Selecting an overly small block size will yield spurious extremes, leading to incorrect parameter estimates; conversely, an excessive block size will limit the amount of extreme data available for analysis, thereby increasing parameter-estimate variability and yielding unstable results. Typically, researchers select block sizes based on a calendar convention, e.g., weekly or monthly blocks. While using calendar conventions is convenient and often yields sufficient results in empirical research, it may not capture the inherent statistical properties of financial return distributions. Financial return distributions exhibit both time-dependent behaviour and heavy tails; therefore, selecting block sizes based solely on calendar conventions can lead to inconsistent estimates.

To address this limitation, this study extends the eight procedures proposed by [30] and introduces a data-driven procedure to select the optimal block size for the GEV distribution. The proposed procedure integrates independence diagnostics, goodness-of-fit testing, and predictive performance evaluation in order to identify a block size that yields stable and reliable extreme value estimates. Our data-driven algorithm begins by extracting block maxima and minima across a range of candidate block sizes. For each block size, the GEV parameters are estimated via maximum likelihood, as in [31]. The adequacy of the fitted distributions is subsequently evaluated using the Anderson Darling, Kolmogorov Smirnov, and Cramér von Mises goodness-of-fit tests. These tests are selected because they provide complementary measures of distributional fit, with the Anderson–Darling statistic placing greater emphasis on the tails of the distribution, which is particularly important in extreme value analysis. Finally, the predictive performance of each fitted model is assessed using out-of-sample error metrics, including the mean squared error (MSE), mean absolute error (MAE), and root mean squared error (RMSE). The complete procedure is summarised in Algorithm 1.

Algorithm 1: Data-Driven Optimal Block Size Selection for GEV

Input: Return series R_t , $t = 1, \dots, T$; candidate block sizes $K = \{5, 6, \dots, 60\}$

Output: Optimal block size k^* and estimated GEV parameters (μ, σ, ξ)

Step 1: Data Partitioning and Scaling;

Compute log-returns R_t series. Partition the data into training (R_{train}) and validation (R_{val}) sets.

Step 2: Candidate Iteration;

For each $k \in K$, repeat Steps 3 through 7:

Step 3: Construct Block Maxima;

Partition R_{train} into $N = \lfloor T_{train}/k \rfloor$ non-overlapping blocks. Extract the maxima for each block:

$$M_{i,k} = \max(R_{(i-1)k+1}, \dots, R_{ik}) \quad \text{for } i = 1, \dots, N$$

Step 4: Independence Diagnostic;

Compute the first-order autocorrelation of the maxima series M_k . Discard k if $|ACF(1)| > 0.2$ to ensure the i.i.d. assumption holds.

Step 5: GEV Parameter Estimation;

Fit the Generalised Extreme Value (GEV) distribution to M_k using Maximum Likelihood Estimation (MLE) to obtain parameters (μ_k, σ_k, ξ_k) .

Step 6: Goodness-of-Fit Validation;

Perform Kolmogorov–Smirnov (KS) and Cramér–von Mises (CvM) tests. Discard k if the p -value for either test is < 0.05 .

Step 7: Out-of-Sample Predictive Performance;

Using the validation set R_{val} , compute the block maxima $M_{val,k}$ and the predicted mean $\hat{E}[M_k]$ from the fitted GEV. Calculate:

$$MSE_k = \frac{1}{n} \sum (M_{val} - \hat{E}[M])^2, \quad MAE_k = \frac{1}{n} \sum |M_{val} - \hat{E}[M]|, \quad RMSE_k = \sqrt{MSE_k}$$

Step 8: Optimal Block Size Selection;

Compute the composite performance score S_k :

$$S_k = \frac{1}{3}(MSE_k + MAE_k + RMSE_k)$$

The optimal block size k^* is selected as:

$$k^* = \arg \min_{k \in K} S_k$$

return k^* and the corresponding parameters (μ, σ, ξ) .

Following the work of [29], the selected optimal block size is then used to estimate the Generalised Extreme Value distribution. The cumulative distribution function of the GEV distribution is defined as,

$$\text{GEV}(x, \xi, \mu, \sigma) = \begin{cases} \exp \left\{ - \left(1 + \xi \left(\frac{x-\mu}{\sigma} \right) \right)^{-\frac{1}{\xi}} \right\}, & \xi \neq 0, \\ \exp \left\{ - \exp \left[- \left(\frac{x-\mu}{\sigma} \right) \right] \right\}, & \xi = 0. \end{cases} \quad (1)$$

The parameter ξ determines the tail behaviour of the distribution. Specifically, $\xi < 0$ corresponds to the Weibull distribution (Type III), $\xi = 0$ corresponds to the Gumbel distribution (Type I), and $\xi > 0$ corresponds to the Fréchet distribution (Type II). For further discussion of these distributions, the reader is referred to [32, 33].

2.3. Generalised Pareto Distribution and POT

In terms of the peak-over-threshold method, we define the excess distribution function, $F_u(x) = P(X - u \leq x | X > u)$ for a random variable, X , where x defines the size of the excesses above the chosen threshold, u . It has been shown in [34, 35] that selecting an appropriate threshold prior to using a Generalised Pareto Distribution is necessary, and a higher threshold yields less biased results; however, it also increases the variance of the selected distribution. Conversely, a lower threshold may yield less variable results; however, it will also increase bias. Therefore, selecting a threshold that balances bias and variance is one of the major challenges in applying the GPD in practice. To address this problem, we present the Stability-Prediction Threshold Optimisation Algorithm (SPTOA), which identifies the optimal threshold based on both the stability of the parameters used in the distribution and the distribution's predictive performance. Our algorithm, i.e., SPTOA, assesses potential thresholds at high quantiles of the return distribution, then selects the threshold that yields the lowest possible value of a combined objective function. This algorithm is defined in Algorithm 2.

Algorithm 2: Stability Prediction Threshold Optimisation Algorithm (SPTOA)

Input: Return series X_1, X_2, \dots, X_T , candidate thresholds $U = \{u_q : q \in [0.90, \dots, 0.99]\}$

Output: Optimal threshold u^*

For each $u \in U$, define exceedances $Y_i = X_i - u$ where $X_i > u$.

Estimate the GPD parameters (σ, ξ) using maximum likelihood estimation.

Compute the stability index

$$S(u) = \left| \frac{d\hat{\sigma}(u)}{du} \right| + \left| \frac{d\hat{\xi}(u)}{du} \right|.$$

Evaluate predictive accuracy using the root mean squared error

$$RMSE(u) = \sqrt{\frac{1}{n_u} \sum_{i=1}^{n_u} (Y_i - \hat{Y}_i)^2}.$$

Define the composite optimisation criterion

$$C(u) = \lambda S(u) + (1 - \lambda) RMSE(u),$$

where $\lambda = 0.5$ assigns equal weight to parameter stability and predictive accuracy.

Select the optimal threshold

$$u^* = \arg \min_{u \in U} C(u).$$

return u^* .

The selected optimal threshold is then used to estimate the Generalised Pareto distribution. The probability density function of the GPD distribution is given by, The probability density function (PDF) is defined as follows: Its cumulative distribution function is defined by:

$$GPD(x, \xi, \mu, \sigma) = \begin{cases} 1 - \left(1 + \xi \left(\frac{x - \mu}{\sigma}\right)\right)^{-\frac{1}{\xi}}, & \text{for } \xi \neq 0, \\ 1 - \exp\left(-\frac{x - \mu}{\sigma}\right), & \text{for } \xi = 0 \end{cases} \quad (2)$$

where x are the returns series $x - \mu \geq 0$ for $\xi \geq 0$, $0 \leq x - \mu \leq -\frac{\sigma}{\xi}$, and $\xi < 0$ after $0 \leq x - \mu \leq -\sigma/\xi$. ξ , is the extreme value index, and σ is the scale parameter. The value of ξ measures the heaviness of the tail, with bigger positive values ($\xi > 0$) indicating a heavy tail. When ($\xi < 0$) is negative, the tail is short (bounded). $\xi = 0$ indicates a light tail.

2.4. Parameter Estimation

The Maximum Likelihood Estimation (MLE) method is employed to estimate the parameters of both the GEV and GPD models. According to [17], the parameters of extreme value distributions can be simultaneously estimated using maximum likelihood methods due to their desirable asymptotic properties.

Given n observations x_1, x_2, \dots, x_n , the log-likelihood function for the GEV distribution is

$$L_{ged} = \mathcal{L}(\mu, \sigma, \xi) = -n \log \sigma - \sum_{i=1}^n \left(1 + \xi \frac{x_i - \mu}{\sigma}\right)^{-1/\xi} - \sum_{i=1}^n \log \left(1 + \xi \frac{x_i - \mu}{\sigma}\right) \quad (3)$$

subject to the constraint

$$1 + \xi \frac{x_i - \mu}{\sigma} > 0.$$

Similarly, after selecting the threshold in the POT framework, the GPD parameters are estimated by maximising the log-likelihood function [36]

$$L_{gpd} = \mathcal{L}(\xi, \sigma) = -n \log \sigma - \left(1 + \frac{1}{\xi}\right) \sum_{i=1}^n \log \left(1 + \xi \frac{x_i}{\sigma}\right). \quad (4)$$

2.5. Model Validation and Selection

To evaluate the adequacy of the fitted models, several goodness-of-fit statistics and model selection criteria are employed. The Cramér–von Mises (CVM), Anderson–Darling (A–D), and Kolmogorov–Smirnov (K–S) tests are used to assess the goodness-of-fit of both distributions [37]. These tests compare the empirical distribution with the theoretical distribution and provide complementary information regarding the quality of fit.

2.6. Risk measures

For a small tail probability, p , and sample size, n , the formula for estimating the Value at Risk (VaR) of returns using a GPD with optimal estimates $(\hat{\sigma}, \hat{\xi})$, threshold, u , and N_u the number of exceedances, is given by:

$$\widehat{\text{VaR}}_p = \begin{cases} u + \frac{\hat{\sigma}}{\hat{\xi}} \left\{ \left(\frac{n}{N_u} p\right)^{-\hat{\xi}} - 1 \right\} & \hat{\xi} \neq 0 \\ u - \hat{\sigma} \log \left(\frac{n}{N_u} (1-p)\right) & \hat{\xi} = 0 \end{cases} \quad (5)$$

where $\hat{\sigma}$ and $\hat{\xi}$ are the maximum likelihood estimates of the GPD parameters, and the expected shortfall is then given as:

$$ES_p = \frac{\text{VaR}_p}{1 - \hat{\xi}} + \frac{\hat{\sigma} - \hat{\xi}u}{1 - \hat{\xi}} \quad (6)$$

3. Model Backtesting and Evaluation

To evaluate the predictive accuracy and robustness of the Value-at-Risk (VaR) and Expected Shortfall (ES) forecasts, this study employs a suite of backtesting procedures. These tests assess both the frequency and magnitude of tail losses, ensuring the models are well-calibrated for extreme market events.

3.1. Tail Risk (TR) Test

Tail risk represents the probability of extreme losses occurring in the left tail of the return distribution. Following [38], we utilise a tail risk statistic that addresses the limitations of standard independence tests by accounting for the magnitude of violations. The TR statistic is defined as:

$$TR = \frac{1}{N} \sum_{t=1}^N \frac{L_t - \text{VaR}_t}{\text{ES}_t - \text{VaR}_t} \cdot I(L_t > \text{VaR}_t) \quad (7)$$

where L_t represents the observed loss and $I(\cdot)$ is an indicator function for VaR breaches.

3.2. Dynamic Quantile (DQ) Test

The Dynamic Quantile test, proposed by [39], evaluates whether VaR violations are independent and identically distributed (*i.i.d.*). The test accounts for temporal interdependence by examining the joint hypothesis that the "hit" variables (violations) are uncorrelated with their own lags. The test is conducted via a Wald statistic to assess the joint nullity of coefficients in the linear regression:

$$\text{Hit}_t = \delta + \sum_{i=1}^p \beta_i \text{Hit}_{t-i} + \sum_{j=1}^q \gamma_j X_j + \epsilon_t \quad (8)$$

3.3. Kupiec Unconditional Coverage Test

To validate the adequacy of the GPD for estimating VaR, we employ the unconditional coverage test developed by [40]. This test is predicated on the principle that the proportion of VaR failures should closely align with the corresponding tail probability level p . A failure occurs when the actual return exceeds VaR for upside risk ($y_t > \text{VaR}_p$) or falls below VaR for downside risk ($y_t < \text{VaR}_p$).

Let x represent the number of violations over a sample of size N . The null hypothesis is $H_0 : E[x/N] = p$. Under H_0 , the likelihood ratio test is defined as:

$$LR_{UC} = -2 \ln \left(\frac{p^x (1-p)^{N-x}}{(x/N)^x (1-x/N)^{N-x}} \right) \sim \chi^2 \quad (9)$$

If the p -value associated with the LR_{UC} statistic is less than the significance level, we reject the null hypothesis and conclude that the model is not appropriate.

3.4. Christoffersen Independence Test

The [41] test examines the independence property—whether the likelihood of a VaR violation on a particular day relies on the result of the previous day. This accounts for the volatility clustering frequently observed in financial time series. Following [42], the combined statistic for conditional coverage (LR_{cc}) follows an asymptotic chi-square distribution:

$$LR_{cc} = LR_{UC} + LR_{ind} \sim \chi^2 \quad (10)$$

3.5. Fissler–Ziegel (FZ) Scoring Rule

To evaluate the joint accuracy of ES and VaR, we utilise the Fissler–Ziegel (FZ) scoring rule [43]. This addresses the elicibility issue of ES by integrating it with VaR within a unified scoring framework:

$$S(v, e, y) = (I - \alpha)(G_1(v) - G_1(y)) + G_2(e) \left(e - v + \frac{1}{\alpha} I(y \leq v)(v - y) \right) - G_2(e) \quad (11)$$

3.6. Joint Quantile–Expected Shortfall Regression Test

[44] proposed a regression-based test to jointly validate VaR and ES by assessing tail frequency and average loss size. This model-agnostic test evaluates conditionally unbiased and dynamically well-calibrated predictions. The full joint null hypothesis is tested using Wald-type statistics to ensure that both the frequency of violations and the mean loss conditional on those violations are consistent with the predicted risk measures.

4. Results and discussion

Table 2 presents descriptive statistics that reveal significant differences in volatility and risk characteristics among the BTC/USD, ZAR/USD, and BWP/USD exchange rate returns. The BTC/USD pair demonstrates the highest standard deviation at 0.0359, significantly surpassing the ZAR/USD at 0.0128 and the BWP/USD at 0.0057, indicating its considerable volatility.

Table 2. Descriptive statistics of exchange rate price returns.

	Observations	Mean	Std.Dev	Maximum	Minimum	Skewness	Kurtosis
BTC/USD	2621	0.00152	0.03586	0.22512	-0.46473	-0.74343	11.79977
ZAR/USD	2621	0.00015	0.01279	0.20088	-0.20411	-0.07458	91.47248
BWP/USD	2621	-0.000123	0.00565	0.05697	-0.05488	-0.16514	12.03358
Test for normality, autocorrelation, and heteroscedasticity							
TEST	BTC/USD		ZAR/USD		BWP/USD		
	Statistic	p-value	Statistic	p-value	Statistic	p-value	
Jarque-Bera	14236.14	0.0001	141.60	0.0001	4587.20	0.0001	
Ljung-Box	66.27	0.0506	63.42	0.0807	128.35	0.0001	
ARCH LM	83.41	0.0016	141.56	0.0001	242.54	0.0001	
Test for unit root and stationarity							
Unit Root	Stat	CV(5%)	Stat	CV(5%)	Stat	CV(5%)	
ADF	-34.79	-1.95	-35.32	-1.95	-38.59	-1.95	
PP	-50.68	-2.86	-48.63	-2.86	-57.05	-2.86	
QLR	28.8	13.42	15.9	13.42	32.7	13.42	

4.1. Preliminary Tail Analysis and EVT Modelling

The main objective of this section is to introduce the extreme value models used in this study, namely the Generalised Extreme Value distribution and Generalised Pareto Distribution. In addition, Figure 1 presents the block maxima and minima of the three currency pairs as initial evidence of asymmetry in tail behaviour. Given the structural differences between extreme losses and extreme gains, it is natural to model the lower tail with the GEV, while the upper tail is better captured by the GPD, allowing for a more flexible characterisation of rare, high-impact events. The figure illustrates the Block Maxima approach, which underpins the application of the GEV distribution within Extreme Value Theory (EVT). Using Algorithm 1, the BTC, ZAR, and BWP return series are divided into non-overlapping temporal blocks, separated by vertical dashed lines, to construct approximately independent and identically distributed (i.i.d.) observations from inherently dependent financial time series. Within each block, only the maximum positive return (circles) and minimum negative return (crosses) are retained, while intra-block observations are suppressed through shading, thereby isolating tail behaviour from central market fluctuations.

A clear heterogeneity in volatility is observed across the three assets. Bitcoin exhibits the widest dispersion between block maxima and minima, reflecting its highly volatile, heavy-tailed nature. In contrast, BWP displays comparatively compressed extremes, indicating lower tail risk, while ZAR lies between the two but still exhibits substantial variability consistent with emerging-market sensitivity. The figure further reveals pronounced tail asymmetry in BTC and ZAR, with negative extremes more pronounced than positive ones, indicating left-skewed

tail behaviour consistent with crash risk and helping explain weaker goodness-of-fit performance in the lower tail under the GEV specification. In addition, volatility clustering is evident, particularly in ZAR, where extreme observations become more concentrated toward the latter part of the sample, potentially reflecting macroeconomic or structural shocks. Overall, by transforming the original return series into a sequence of block-wise extremes, the methodology satisfies the conditions of the Fisher–Tippett–Gnedenko theorem, thereby justifying the use of the GEV distribution and enabling a more robust analysis of tail risk, as well as improved estimation of extreme quantiles and return levels in financial markets.

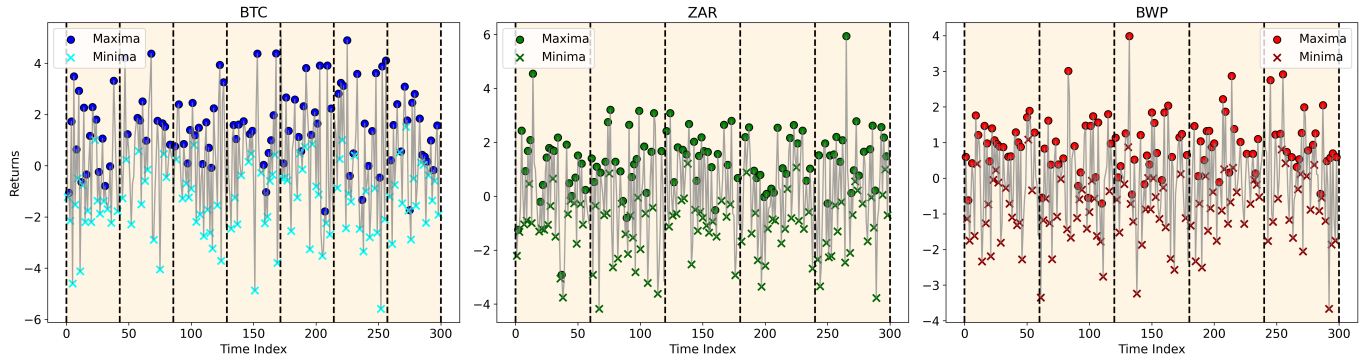


Figure 1. Block maxima and Minima

4.1.1. Block Maxima Analysis and Parameter Estimation To evaluate the block maxima of gains and minima of losses in returns on BTC / USD, ZAR / USD, and BWP / USD, we then apply our proposed Algorithm 1. As a result of this procedure, a total of 553 block maxima and minima are available for the weekly return series. Then, the largest and smallest values for each week are selected for subsequent analysis. The block maxima approach is utilised to select weekly extremes from the daily returns of BTC/USD, ZAR/USD, and BWP/USD. The maximum and minimum are fitted separately to the Generalised Extreme Value Distribution. In addition to providing the maximum likelihood estimates of the GEVD weekly maxima and minima for BTC / USD, ZAR / USD, and BWP / USD, along with their respective standard errors as reported in Table 3, these parameters serve as important tools in the prediction of return levels and the calculation of their corresponding 95% confidence intervals as emphasised by [45]. The Shape Parameter $\hat{\xi}$ is positive for all three currencies. It is counterintuitive that the estimation of $\hat{\xi}$ as positive appears to contradict the findings of [15], which models monthly returns of ZAR / USD and determines that there is a statistically significant negative shape parameter. The results indicate that the exchange rates have fat tails and thus belong to a family of Fréchet distributions [32]. On both sides of the curve, based on results for minima and maxima, the Extreme Value Indices for BTC / USD were found to be larger than those for ZAR / USD. These results agree with [25] in their study of monthly returns for BTC/USD and ZAR/USD. Therefore, because the associated EVIs for BTC/USD represent both short- and long-term trading positions, and are higher than those for ZAR / USD, BTC / USD has been shown to be riskier than ZAR / USD. BTC / USD has an EVI of 0.267 with a standard error of 0.035, representing the probability of losing money (losses), whereas it has an EVI of 0.129 with a standard error of 0.037, representing the probability of gaining money (gains). The difference between these two statistics implies that Bitcoin is more prone to large downward swings than upward ones. Conversely, the probability of large upward swings in ZAR / USD (gains) is greater than large downward swings (losses). Since the shape parameters estimated for BWP / USD are 0.025 for losses and 0.017 for gains, however, since the 95 per cent confidence intervals for these estimates contain zero, they are not statistically significant. Thus, there is very little evidence supporting the idea that BWP has fat tails and therefore may be modelled using a class of Gumbel distributions. This conclusion also supports the hypothesis that the BWP is perceived as a more stable currency.

Table 3. Generalised Maximum Likelihood Estimates for the GEV Distribution

Returns	Minima						Maxima					
	$\hat{\xi}$	$se\hat{\xi}$	$\hat{\sigma}$	$se\hat{\sigma}$	$\hat{\mu}$	$se\hat{\mu}$	$\hat{\xi}$	$se\hat{\xi}$	$\hat{\sigma}$	$se\hat{\sigma}$	$\hat{\mu}$	$se\hat{\mu}$
BTC/USD	0.267	0.035	0.020	0.001	0.024	0.001	0.129	0.037	0.022	0.001	0.029	0.001
ZAR/USD	0.080	0.018	0.005	0.000	0.007	0.000	0.113	0.022	0.006	0.000	0.008	0.000
BWP/USD	0.025	0.019	0.004	0.000	0.004	0.000	0.017	0.017	0.003	0.000	0.004	0.000

The GEV estimates reveal distinct differences in the tail behaviour of the three exchange rate return series, as indicated by the shape parameter. Bitcoin BTC/USD exhibits the largest positive ξ for losses 0.267, confirming heavy-tailed dynamics and highlighting its high susceptibility to extreme downside realisations, which is consistent with the findings of [16]. A critical finding in this study is the divergence in shape parameters for Bitcoin, with a negative parameter for gains, and a positive parameter for losses, $\hat{\xi} = 0.260$, revealing a complex, asymmetric risk profile. This suggests that while upward price movements are constrained, likely due to periodic profit-taking and regulatory ceilings, downward price crashes are theoretically unbounded and heavy-tailed. Such crash-o-phobia in the cryptocurrency market is frequently triggered by sudden regulatory announcements or shifts in social media sentiment, initiating feedback loops of panic-selling that far exceed the magnitude of any FOMO-driven rallies. These results reinforce the view that decentralised digital assets introduce a fundamentally different risk structure in financial markets, in which the potential for extreme depreciation is not mirrored by an equivalent potential for extreme appreciation.

Figure 2 displays diagnostic plots for the models applied to BTC/USD maxima (left) and minima (right). In Figure 2a, we present the graphical goodness-of-fit plots for BTC/USD returns gains. The probability and quantile plots exhibit near-linear characteristics, indicating a strong fit. The return levels fall within the expected confidence bands. Additionally, the density plot serves as a reliable estimate of the data histogram. Overall, the model demonstrates a commendable fit to the BTC/USD gains data but struggles with the most extreme gains, showing fatter right tails than predicted. In Figure 2b, the probability and quantile plots for the BTC/USD returns closely align with the straight line, indicating that the GEV model is well-suited for BTC/USD losses. The return level plots show no significant deviations from the specified confidence level band. Additionally, the density plot provides an accurate estimate of the histogram of the BTC/USD losses data but underestimates extreme losses, indicating fat left tails. For BTC/USD, the diagnostics show that although the models effectively capture the overall structure of extremes, sharp losses (minima) are less predictable and exhibit heavier tails compared to gains (maxima).

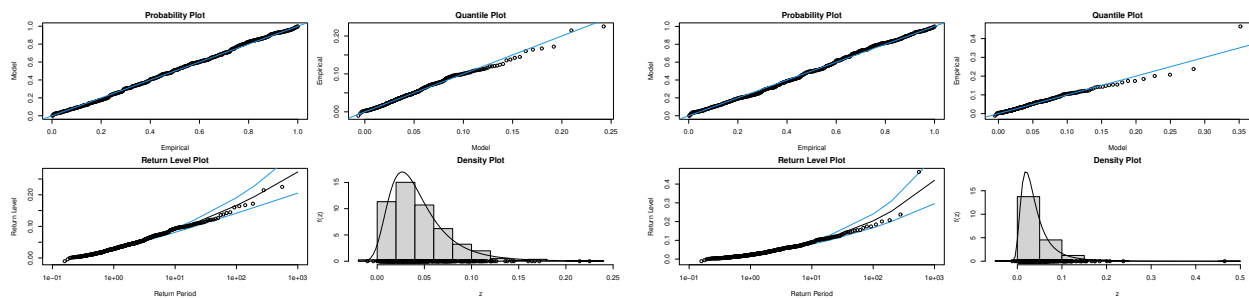


Figure 2. Model diagnostics for the maxima and minima of the BTC/USD returns

Figure 3 displays diagnostic plots for the models applied to ZAR/USD maxima (left) and minima (right). In Figure 3a, we show the graphical goodness-of-fit plots for BTC/USD return gains. The probability and quantile plots exhibit near-linear characteristics, suggesting a strong fit, with return levels falling within the expected confidence bands. The density plot also reliably estimates the data histogram, indicating that the model fits the ZAR/USD gains data well. In Figure 3b, the probability and quantile plots for BTC/USD return losses closely

align with the straight line, demonstrating that the GEVD model is well-suited for these losses. The return level plots show no significant deviations from the specified confidence level band, and the density plot accurately estimates the histogram of the ZAR/USD losses data. The ZAR/USD exchange rate may exhibit greater left tail risks, indicating a higher likelihood of sudden depreciation, than right tail risks. This aligns with financial intuition, as currencies typically depreciate more abruptly and unpredictably than they appreciate.

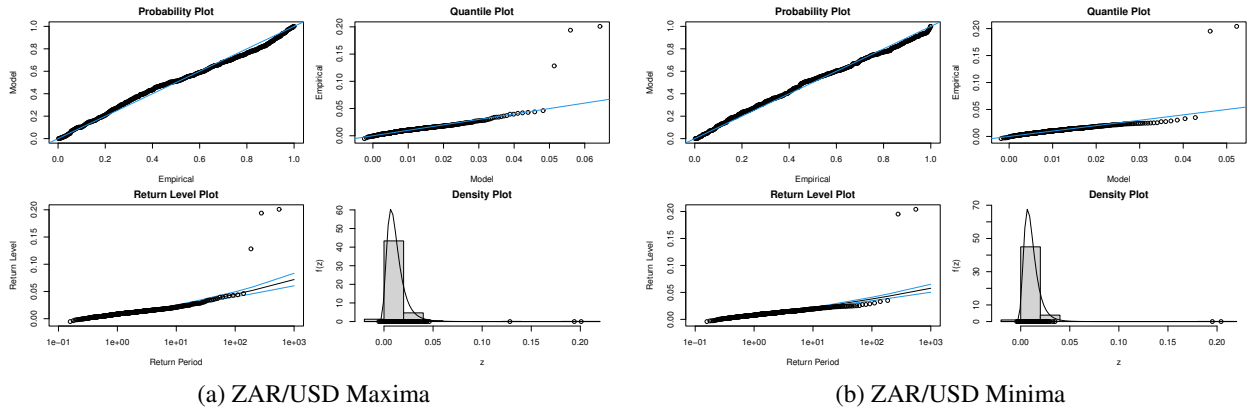


Figure 3. Model diagnostics for the maxima and minima of the ZAR/USD returns

Figure 4 presents diagnostic plots for the models applied to BWP/USD maxima (left) and minima (right). In Figure 4a, we illustrate the graphical goodness-of-fit plots for BWP/USD return gains. The probability and quantile plots align closely with the diagonal, with only minor deviations observed in the upper tail, indicating a strong fit. The return level and density plots confirm that moderate positive extremes are well-represented, though very large surges are slightly underestimated. The EVT diagnostics suggest that the fitted models adequately capture the maxima of the BWP/USD exchange rate. In Figure 4b, the minima diagnostics indicate more pronounced deviations in the lower quantiles, wider confidence intervals in the return level plot, and heavier left tails in the density plot. This suggests that extreme depreciations are less predictable than appreciations. Overall, the results highlight asymmetric tail behaviour, with the BWP/USD exhibiting fatter left tails, indicating a higher likelihood and greater severity of extreme losses than extreme gains.

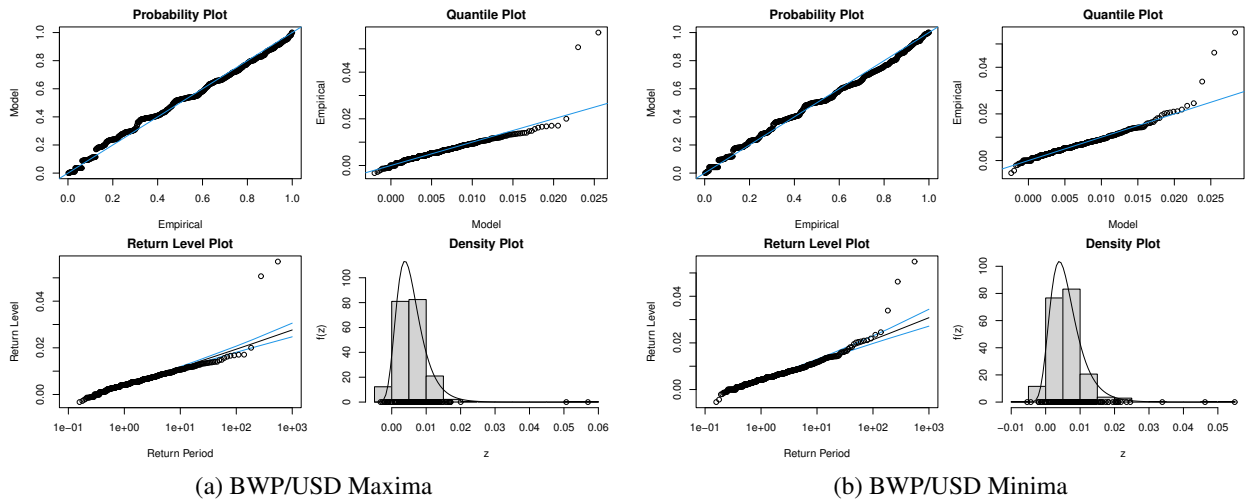


Figure 4. Model diagnostics for the maxima and minima of the BWP/USD returns

The Generalised Extreme Value estimates reveal distinct differences in the tail behaviour of the three exchange rate return series, as indicated by the shape parameter ξ . Bitcoin (BTC/USD) exhibits the largest positive ξ , confirming heavy-tailed dynamics and highlighting its high susceptibility to extreme return realisations. This firmly establishes Bitcoin as the riskiest of the assets under consideration, with conventional volatility-based metrics insufficient to capture its propensity for rare but severe price movements. It is worth noting that these are in line with those reported in [16], who found that BTC/USD is riskier than the ZAR/USD series they used. The South African Rand also yields a positive, though smaller, ξ , suggesting moderately heavy tails. While less extreme than Bitcoin, the Rand remains exposed to tail risks stemming from global commodity market shocks and shifts in investor sentiment in emerging markets. In contrast, the Botswana Pula (BWP/USD) exhibits a comparatively smaller ξ , which may be close to zero or even negative, implying thinner tails relative to the other series. This positions the Pula as the least risky of the three, reflecting its relatively managed exchange-rate regime and a more stable macroeconomic environment. Taken together, the hierarchy of risk implied by the GEV estimates is clear: BTC/USD is the most exposed to extreme downside shocks, followed by ZAR/USD, with BWP/USD the most stable, albeit not entirely insulated from occasional sharp depreciations.

4.1.2. Generalised Pareto Distribution and Peaks over Threshold The first and most critical step in GPD modelling is identifying the optimal threshold u . To avoid subjectivity, we therefore apply our data-driven Algorithm 2 and evaluate candidate thresholds between the 90th and 99th percentiles of the weekly return series. Following [37], we identify thresholds where the model parameters, specifically the shape parameter ξ and the modified scale parameter, demonstrate stability. Optimal thresholds for the gains (maxima) of BTC/USD, ZAR/USD, and BWP/USD are identified at 3.9728, 1.2934 and 0.7971, respectively. For the loss series (minima), the optimal thresholds are -3.7520, -1.1575 and -0.7178. These selections are justified by the parameter stability plots and cost function, which exhibit clear regions of linearity at these levels, as shown in Figure 5. While Figure 6 presents both extremes above the selected threshold.

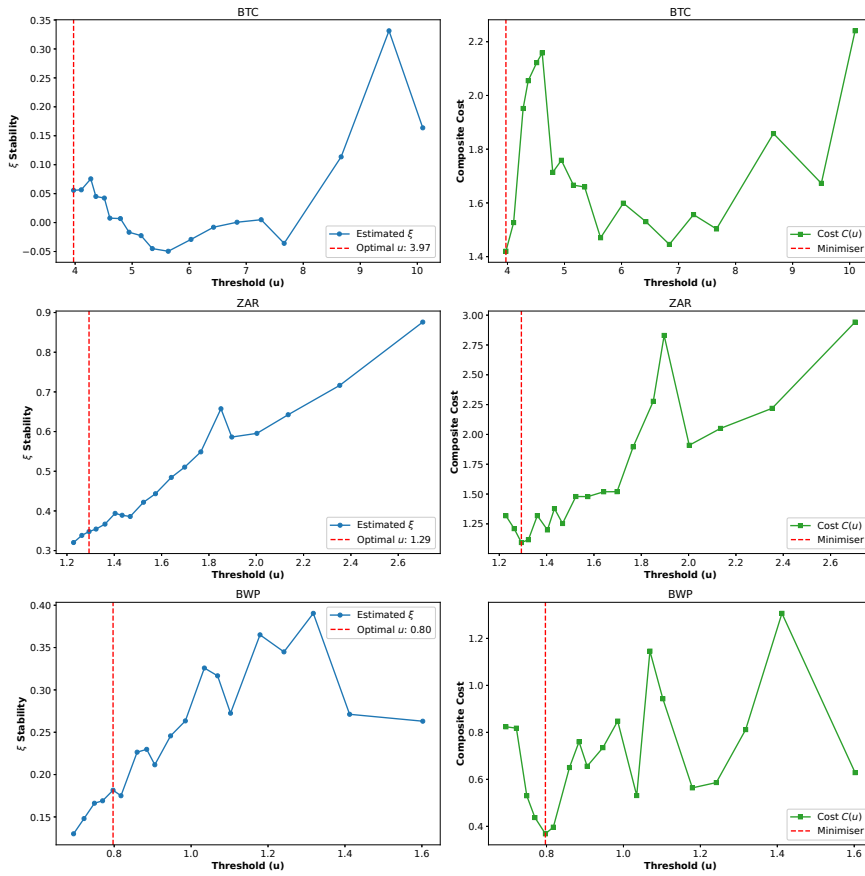


Figure 5. Stability plot for the shape parameter of the GPD

Once the thresholds are established, the GPD parameters are estimated using Maximum Likelihood Estimation. This method provides efficient estimates of the scale parameter σ , which governs the dispersion of exceedances, and the shape parameter ξ , which determines the tail behaviour. Table 4 summarises these estimates along with their standard errors and 95% confidence intervals.

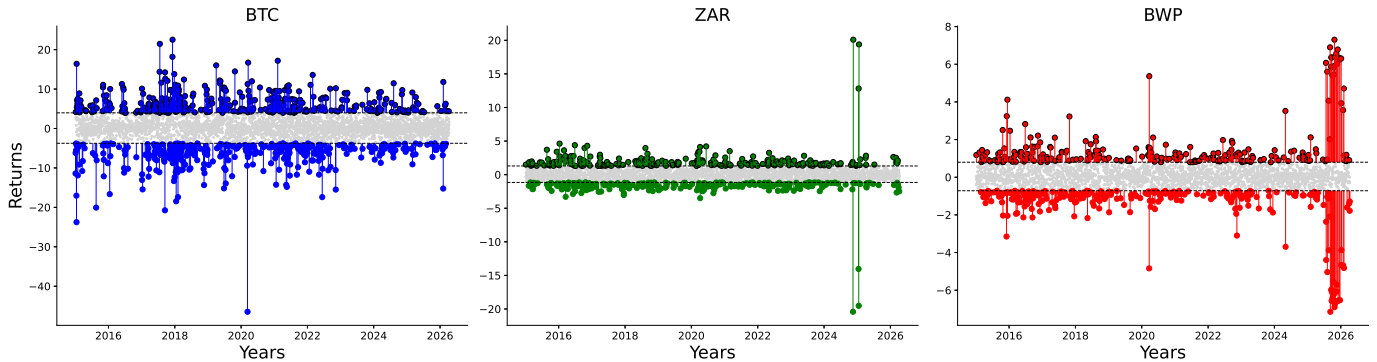


Figure 6. Exceedances high or below selected thresholds

In addition to the heterogeneity in tail behavior shown in Table 4, the Generalized Pareto Distribution (GPD) model also shows that there are significant variations among the shape parameter estimations obtained from

the Generalized Pareto Distribution (GPD) for each of Bitcoin (BTC), the South African Rand (ZAR), and the Botswana Pula (BWP). For BTC, both tails have a relatively low positive shape parameter. Specifically, the shape parameter for the lower tail has an estimate of $\hat{\xi} = 0.1138$, with a standard error of 0.0539. Similarly, for the upper tail, we obtain an estimated value for the shape parameter of $\hat{\xi} = 0.0612$, with a corresponding standard error of 0.0586. Hence, the estimated shape parameters provide evidence of mildly heavy-tailed behaviour in both tails of the distribution, with slightly greater tail risk on the downside. While the two estimated values are similar in size, their respective standard errors are equal. Therefore, although the two estimated values indicate relatively modest levels of tail heaviness, they do so in a statistically reliable manner, given their similarly sized standard errors. Conversely, ZAR displays significantly more tail dependence in both loss and gain directions than BTC. More specifically, its lower tail shape parameter is estimated as $\hat{\xi} = 0.2698$, with a standard error of 0.0578, while the upper tail's shape parameter is estimated as $\hat{\xi} = 0.3280$, with a standard error of 0.0722. Therefore, these values indicate much heavier tails than those of BTC and, therefore, a greater likelihood of extreme exchange rate movements. This type of behaviour is typical of emerging market currencies, which are subject to various macroeconomic shocks and capital flows, resulting in unstable exchange rates and shifts in investor perceptions of risk.

Despite the fact that the standard errors associated with these heavy-tail effects were smaller than the actual sizes of the effect estimates themselves, the reliability of these effects can still be confirmed. The most extreme heavy-tail behaviour is exhibited by BWP; the shape parameters for both its lower and upper tails re found to be equal to approximately $\hat{\xi} = 0.6494$, with a standard error of 0.0971, and $\hat{\xi} = 0.6697$, with a standard error of 0.1019, respectively. As such, these estimates demonstrate extremely heavy-tailed behaviour in both directions and thus a high likelihood of rare but substantial movements relative to BTC and ZAR. Furthermore, because both the lower- and upper-tail estimates are roughly equal, the symmetry of the extremal structure can be implied. Although BWP's estimates are significantly larger than those of BTC and ZAR, the fact that the former's standard errors are somewhat larger than those of BTC and ZAR does not diminish the magnitude of the difference between them. Thus, despite some uncertainty regarding BWP's extreme-value behaviour, reflected in its larger standard errors than those of BTC and ZAR, the fact remains that BWP's extreme-value behaviour is quite different from that of BTC and ZAR. Regardless of whether one considers either tail or overall excesses, it appears that roughly equal numbers of thresholds were selected for each tail across all three assets. This suggests that regardless of which direction one looks, the number of extreme observations captured by selecting a given set of thresholds is roughly equal. However, there are large differences in scale among the three assets; specifically, BTC is farthest from zero in absolute terms, indicating high variability in extreme cryptocurrency returns. On the other hand, ZAR and BWP are closer to zero but appear to exhibit proportionally greater tail sensitivity than BTC.

Overall, our analysis demonstrates a distinct hierarchy in tail heaviness among our three assets: BTC shows mildly heavy tails, while ZAR exhibits moderately strong extreme behaviour. Finally, BWP represents perhaps the greatest tail risk among all three assets. Our findings emphasise the need for asset-specific tail modelling and illustrate that emerging-market currencies – especially BWP – may be more susceptible to extreme-value risk than even highly volatile cryptocurrency markets when considering GPD shape-parameter estimates.

Table 4. Maximum Likelihood Parameter Estimates for GPD

Returns	Minima (Lower Tail)				Maxima (Upper Tail)			
	\hat{u}	N_{exc}	$\hat{\sigma}$ (se)	$\hat{\xi}$ (se)	\hat{u}	N_{exc}	$\hat{\sigma}$ (se)	$\hat{\xi}$ (se)
BTC	-3.7520	356	2.7924 (0.2103)	0.1138 (0.0539)	3.9728	389	2.4748 (0.1916)	0.0612 (0.0586)
ZAR	-1.1575	278	0.4230 (0.0346)	0.2698 (0.0578)	1.2934	248	0.4907 (0.0460)	0.3280 (0.0722)
BWP	-0.7178	263	0.3425 (0.0369)	0.6494 (0.0971)	0.7971	246	0.3524 (0.0396)	0.6697 (0.1019)

Note: \hat{u} is the threshold, N_{exc} is the number of exceedances, and values in parentheses denote standard errors.

In our analysis of the lower-tail and upper-tail as reported in Table 4, generalised Pareto distribution fits to exceedance data from Bitcoin (BTC), South African Rand (ZAR) and Botswana pula (BWP), and we found that BTC shows a very good generalised Pareto distribution fit in terms of both sample behavior and tail behaviour;

however, we do find some variation among the three asset classes in terms of sample size. We also see evidence of non-linearities in the Q-Q-plots. For example, BTC has an almost perfectly linear relationship between observed values and quantiles of the fitted gpd, while there is a clearly visible deviation for the ZAR currency data: specifically, many of its upper tail observations fall well above the diagonal line, which suggests that the fitted model may be underestimating extreme devaluation events. Similarly, BWP shows a distinctly curved pattern followed by a distinct flattening, which could suggest either that the model is misspecified, or simply that the tails of the BWP distribution are far too complex to be modelled accurately using a simple gpd. The survival plots confirm these findings: they show that BTC's empirical survival function matches the theoretical survival function remarkably well up to the 10^{-3} confidence level. In contrast, the zrh and BWP empirical survival functions deviate from their respective theoretical survival functions at higher confidence levels, again implying less reliable models in those two cases. While it appears that the empirical survival curves for all three currencies terminate at a 10^{-3} confidence level, this does not represent a limit on the actual probability of extreme events. Rather, it represents only an artefact of finite sampling. That is, because the GPD is defined to allow extrapolation beyond the bounds of what is actually sampled, it can provide estimates of the likelihood of events far more extreme than anything seen during our observation period. In addition to these overall trends, we note that BTC's estimated return levels closely match their corresponding theoretical levels. However, both ZAR and BWP exhibit much larger variances in their estimated return levels than BTC does. Furthermore, in both of the latter two cases, these variances increase significantly at longer return horizons. Overall, then, we believe that our diagnostics demonstrate that the GPD is capable of providing reliable fits for tail risk associated with cryptocurrencies, but currency markets, particularly BWP, are likely to exhibit more irregular extreme-value behaviours that will require additional structural modifications before accurate modelling is possible.

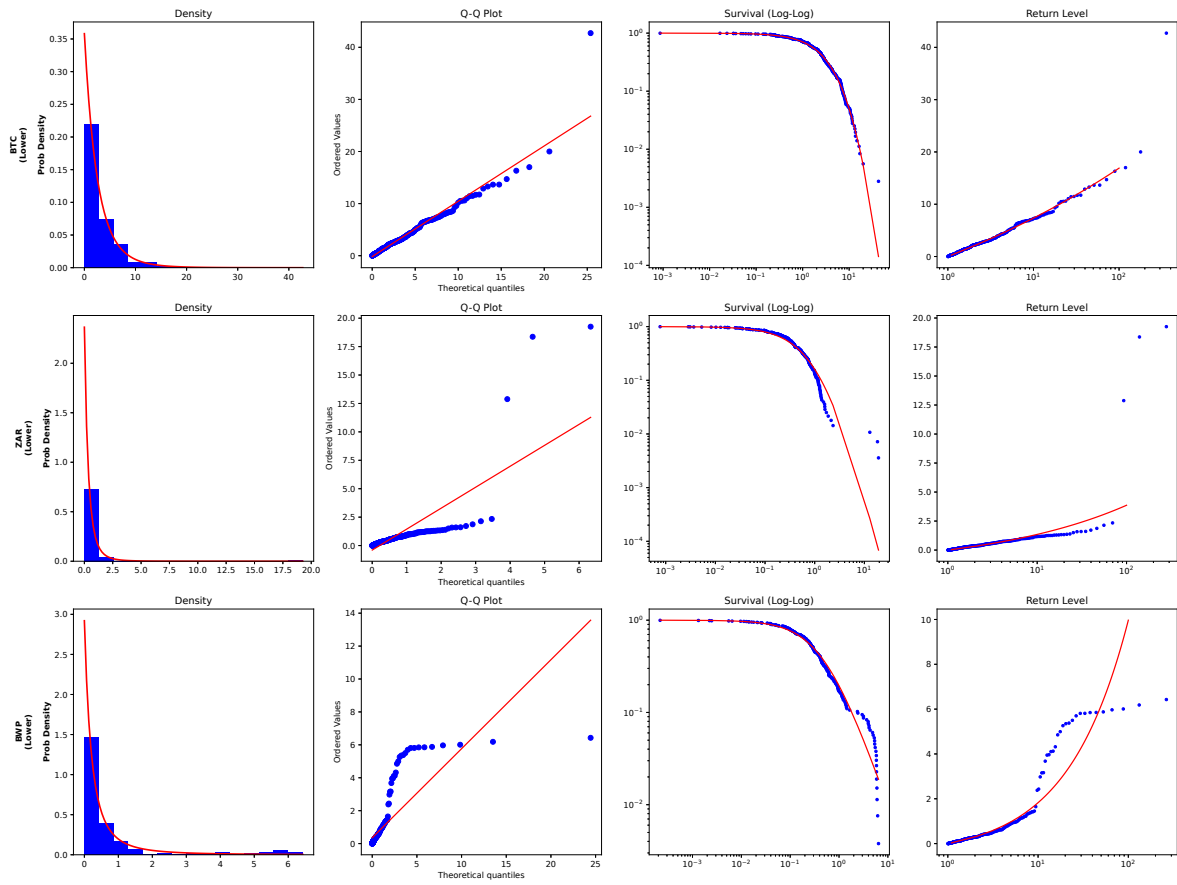


Figure 7. GPD diagnostic plots for BTC/USD, ZAR/USD, and BWP/USD across lower tail exceedances

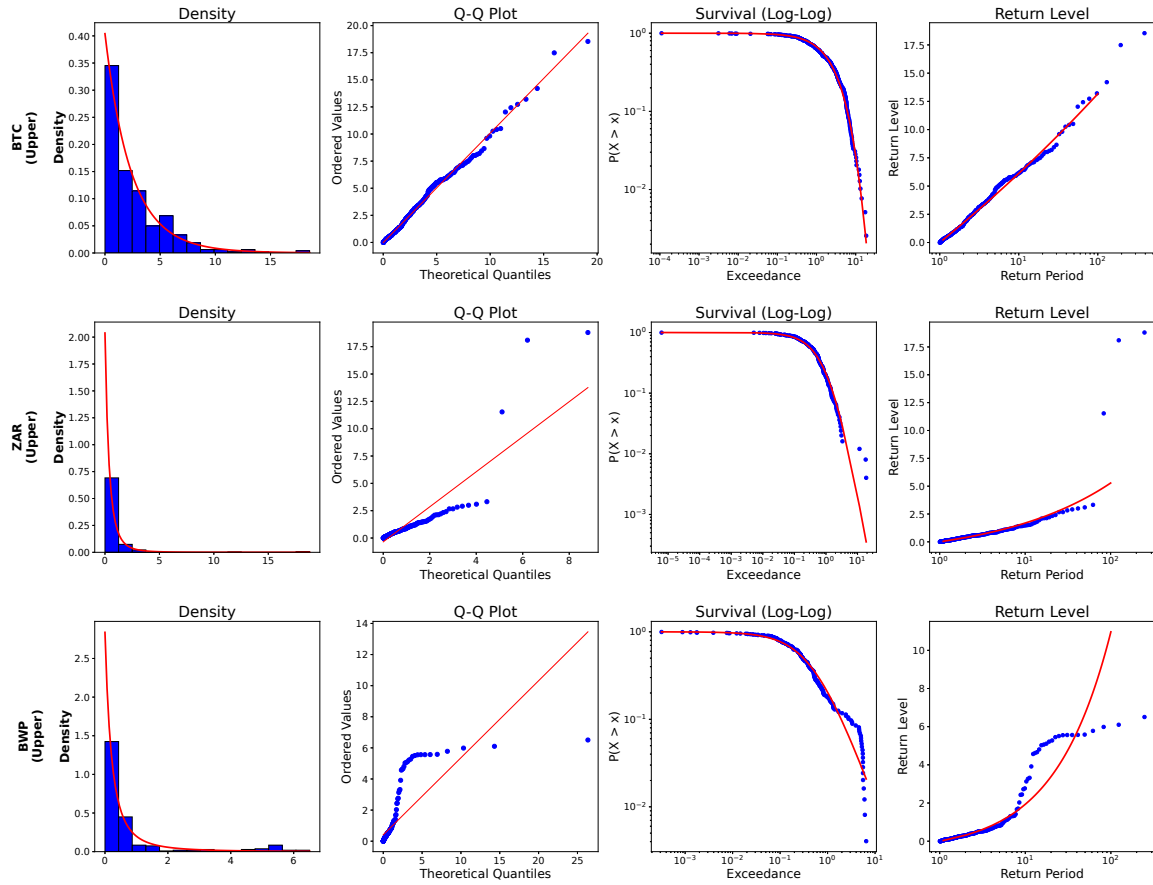


Figure 8. GPD diagnostic plots for BTC/USD, ZAR/USD, and BWP/USD across upper tail exceedances

The goodness-of-fit results presented in Table 5 provide consistent empirical evidence on the relative performance of the Generalised Pareto Distribution (GPD) and Generalised Extreme Value (GEV) models across both tails. For BTC/USD, the GPD demonstrates a strong fit in both the upper and lower tails, with Kolmogorov–Smirnov p-values of 0.3495 and 0.6835, respectively, and relatively low Anderson–Darling statistics, indicating that the model captures extreme behaviour with high accuracy. By contrast, while the GEV provides a similarly good fit for maxima, reflected in a KS p-value of 0.4032 and a CvM p-value of 0.4395, it performs poorly for minima, where the KS p-value falls to 0.0394 and the CvM p-value to 0.0355, accompanied by a large Anderson–Darling statistic, confirming a clear model misfit in the lower tail. A comparable pattern is observed for ZAR/USD, where the GPD maintains robustness in the upper tail, with a KS p-value of 0.6033, and provides an acceptable, though weaker, fit in the lower tail with a KS p-value of 0.1030, whereas the GEV shows only marginal adequacy for maxima and is decisively rejected for minima, as evidenced by a KS p-value of 0.0001 and a CvM p-value of 0.0009, together with a substantially elevated Anderson–Darling statistic. In the case of BWP/USD, both models perform reasonably well across tails, with the GPD exhibiting stable goodness-of-fit and the GEV also producing acceptable KS and CvM results, although higher Anderson–Darling statistics suggest mild tail irregularities under the block maxima.

Table 5. Goodness-of-Fit Results for GPD (POT) and GEV (Block Maxima)

Asset	Tail	Model	KS p-value	CvM p-value	AD Stat
BTC	Upper	GPD	0.3495	0.1585	0.9021
BTC	Upper	GEV	0.4032	0.4395	0.7923
BTC	Lower	GPD	0.6835	0.0813	0.4482
BTC	Lower	GEV	0.0394	0.0355	3.3979
ZAR	Upper	GPD	0.6033	0.1345	0.8652
ZAR	Upper	GEV	0.0830	0.0749	2.5088
ZAR	Lower	GPD	0.1030	0.3486	2.1767
ZAR	Lower	GEV	0.0001	0.0009	7.7640
BWP	Upper	GPD	0.3437	0.1862	1.4904
BWP	Upper	GEV	0.3460	0.4432	4.7273
BWP	Lower	GPD	0.4342	0.1660	1.2451
BWP	Lower	GEV	0.4809	0.2513	4.2588

Note: The critical value of the Anderson–Darling test at the 5% level is 2.492.

The data collected indicate an imbalance in model applicability across asset classes and tails. As expected, the GPD model generally provided a consistent and reliable representation of extreme behaviour. This was especially true in the lower tails of financial distributions, where extreme loss potential is typically highest. On the other hand, the GEV model was only applicable to modelling upper-tail extremes, as evidenced by no rejection of the goodness-of-fit test. Marketwise, this is important because the processes that generate extreme losses in financial markets are distinctly different from those that generate extreme returns. Moreover, the GEV model’s failure in the lower tails of Bitcoin and the South African Rand indicates that it cannot capture the sudden drop-off associated with crashes in these markets. Furthermore, the consistency of the GPD model across all tails, particularly for downside risk, supports the use of the POT methodology in various risk management applications, such as estimating extreme quantiles, conducting stress tests, and identifying low-probability/high-impact events. The GEV (Block Maxima) model provided a good estimate of upper-tail gains; however, it had notable shortcomings in representing the heavy-tailed nature of downside risks for both Bitcoin and the South African Rand. Conversely, the GPD (Peaks Over Thresholds) methodology met or exceeded goodness-of-fit requirements for all assets and tails and therefore appears to be a more dependable method for identifying the asymmetric/extreme risk profiles of both cryptocurrencies and emerging market exchange rate exposures.

4.2. Evaluating Downside and Upside Risk using the GPD

The estimates for Value-at-Risk (VaR) and Expected Shortfall (ES) derived from the Generalised Pareto Distribution (GPD) models are summarised in Tables 6 and 7.

Table 6. VaR estimates for the GPD models.

	BTC/USD		ZAR/USD		BWP/USD	
	Losses	Gains	Losses	Gains	Losses	Gains
90%	0.0980	0.0568	0.0206	0.0238	0.0063	0.0023
95%	0.1199	0.1093	0.0219	0.0270	0.0175	0.0267
99%	0.1888	0.1802	0.0311	0.0479	0.0481	0.0574

Although the estimated shape parameters for each of the three time series are positive, the distinct asymmetry in tail heaviness for BTC/USD is a critical feature of the GPD model. Specifically, the lower-tail shape parameter is nearly double that of the upper-tail shape parameter, at 0.1138 versus 0.0612, indicating a highly complex risk profile in which downward price movements exhibit much heavier tails than upward movements. In addition to this “crash-o-phobia”, the cryptocurrency market also appears to be susceptible to rapid feedback loops of panic selling, driven by negative news events such as regulatory announcements and changes in social media sentiment.

While it is possible for cryptocurrency markets to experience sharp price increases during FOMO-driven buying frenzies, the risk of extreme price appreciation often seems greater than that of extreme price depreciation.

Unlike cryptocurrencies, emerging-market currencies appear to exhibit structural patterns in their behaviour. For example, the upper-tail shape parameter for ZAR/USD is estimated at 0.3280, exceeding the estimated lower-tail shape parameter of 0.2698. Similarly, the Botswana Pula exhibits a relatively high degree of tail symmetry and extreme heaviness, with estimated lower- and upper-tail shape parameters of 0.6494 and 0.6697, respectively. These findings suggest that, in these fiat-currency markets, the risk of extreme price appreciation may exhibit a heavier tail profile than the risk of extreme price depreciation, clearly providing evidence of a fundamental difference in the tail-risk structures of decentralised digital assets versus traditional currency markets.

Table 7. ES estimates for the GPD models.

	BTC/USD		ZAR/USD		BWP/USD	
	Losses	Gains	Losses	Gains	Losses	Gains
90%	0.1369	0.1168	0.0331	0.0443	0.0242	0.0296
95%	0.1665	0.1521	0.0452	0.0635	0.0372	0.0453
99%	0.2597	0.1997	0.1286	0.1870	0.0729	0.0650

These empirical estimates provide concrete, actionable metrics for practical risk management and institutional policymaking, extending beyond theoretical modelling. For a trader or hedge fund managing a BTC/USD position, the 99% ES of 0.2597 implies that, in the event of a threshold breach, the average expected loss is approximately 26% of the total investment. Such pronounced tail risk necessitates strict deleveraging; for example, a trader employing 4 times leverage would face total capital depletion during a single extreme event of this magnitude. Consequently, we recommend a maximum leverage ratio of 2 times for Bitcoin portfolios to maintain resilience against these identified non-linear shocks. From a central bank perspective, the results offer a data-driven blueprint for reserve management. The South African Reserve Bank should consider maintaining a “Tail-Risk Buffer” in its foreign exchange reserves equivalent to at least 19% of its short-term external debt obligations, directly reflecting the ZAR/USD 99% ES of 0.1870. Meanwhile, the stability of the Botswana Pula BWP/USD suggests that the current peg mechanism is effective, although the Bank of Botswana should maintain a reserve policy that is “ZAR-proximal” to allow for intervention during periods of regional currency stress. Economically, these findings highlight the risk of volatility transmission. Recent evidence indicates that the annualised volatility of Bitcoin is approximately four times higher than that of traditional equity markets [46]. While it is sometimes viewed as a store of value, it cannot be considered a low-risk vehicle for wealth preservation relative to fiat currencies, which retain superior short-term stability even in high-inflation environments [47]. Although small allocations of 1–3% may enhance risk-adjusted returns [49, 48], standard GARCH-EVT frameworks often lack the long-run robustness required to accommodate rapid structural shifts [50]. This motivates the development of dynamic hybrid models that integrate deep learning with EVT to better capture both downside and upside risks [51]. Furthermore, the growing adoption of digital assets introduces persistent market pressures that can affect treasury yields and the stability of fiat currencies [52], necessitating stricter prudential controls to mitigate systemic instability arising from volatility spillovers between cryptocurrencies and commodity markets [53]. Overall, these results reinforce the view that cryptocurrencies introduce a fundamentally different risk structure into foreign exchange markets, characterised by heavier tails and an elevated need for continuous systemic risk monitoring.

4.3. Evaluating the Accuracy of Risk Measures

To assess the adequacy of the GPD-based risk estimates, a comprehensive suite of backtesting procedures is applied to both Value-at-Risk (VaR) and Expected Shortfall (ES). While previous studies, such as those by Naradh et al. [54], focus exclusively on backtesting VaR, this study recognises that marginal VaR backtesting is inherently limited in its ability to capture the full spectrum of tail risk. The primary problem with focusing only on VaR backtesting is that VaR is not sub-additive, a fundamental property required for a risk measure to be considered coherent [55]. This lack of sub-additivity means VaR can fail to account for the diversification of risk and, more critically, ignores the severity of violations. A model may pass a VaR backtest by accurately predicting the

frequency of exceptions while failing to capture the magnitude of those losses when they occur. This tail blindness creates a dangerous incentive for risk managers to adopt models that satisfy regulatory capital requirements based on frequency, while remaining exposed to catastrophic losses in the distribution's farthest reaches.

Furthermore, because Expected Shortfall is not elicitable, it cannot be backtested as a standalone measure using traditional point-forecast methods. By adopting a joint backtesting framework, this analysis overcomes the elicibility challenge. Following the work of [56], which represents the first empirical finance application to employ joint backtesting of VaR and ES, this study utilises the mathematical foundations established in [43]. This approach allows for a scoring function that penalises the model not only for failing to predict the occurrence of an extreme event but also for underestimating the average loss following a VaR breach, providing a mathematically coherent validation of the financial risk structure. The results, summarised in Table 8, provide a rigorous statistical validation of the model's predictive performance across a range of confidence levels. The findings indicate that the Generalised Pareto Distribution provides a satisfactory fit for both gains and losses across the BTC/USD, ZAR/USD, and BWP/USD exchange rate series. The model's adequacy is supported at the 95% and 99% confidence levels, as all p-values lie comfortably above the 5% significance threshold. The upper panel of Table 8 evaluates the VaR estimates using the Kupiec and Christoffersen tests. The Kupiec test p-values, ranging from 0.2211 to 0.9368, suggest that the model is well-calibrated and does not systematically under- or overestimate the frequency of extreme losses. This is further corroborated by the Christoffersen test p-values ranging from 0.1294 to 0.6964, indicating no clustering in exceedances and implying that the model adapts effectively to evolving market volatility. Additionally, high p-values from the Dynamic Quantile test confirm that VaR exceedances are not autocorrelated. These findings are consistent with recent evidence where GPD-based models effectively capture extreme market movements [57, 58], and adapt to time-varying tail dynamics [59].

Table 8. Backtesting m-ahead of Value-at-Risk and Expected Shortfall

Panel A: Value-at-Risk Backtesting					
Proportion	Violations	TR Test	DQ Test	Christoffersen	Kupiec
95%	7	0.5964	0.9961	0.5158	0.7871
95%	13	0.5870	0.8964	0.6964	0.9368
95%	10	0.7904	0.5064	0.1294	0.9221
95%	19	0.1964	0.8964	0.5764	0.8581
99%	18	0.9552	0.9785	0.1313	0.3337
99%	16	0.0937	0.0508	0.4963	0.8741
99%	12	0.8897	0.8963	0.1879	0.2211
99%	19	0.8897	0.8963	0.1879	0.3587
Panel B: Expected Shortfall Backtesting					
Proportion	Violations	TR Test	DQ Test	FZ Test	Joint Quantile
95%	N/A	0.3885	0.3049	0.2146	0.2879
95%	N/A	0.3991	0.5281	0.3047	0.8987
95%	N/A	0.9955	0.5139	0.1306	0.7871
95%	N/A	0.8772	0.0596	0.1433	0.5989
99%	N/A	0.7768	0.7872	0.7287	0.7898
99%	N/A	0.9508	0.9742	0.1879	0.1121
99%	N/A	0.8897	0.7872	0.2320	0.0975
99%	N/A	0.0232	0.2011	0.1872	0.1587

The lower panel of Table 8 reports the results for Expected Shortfall, where both the Fissler–Ziegel (FZ) and Joint Quantile tests produce p-values consistently above the 0.05 threshold. The FZ results range from 0.1306 to 0.7287, indicating that the ES forecasts are statistically consistent with realised losses beyond the VaR threshold. These joint forecasts are reliable under strictly consistent scoring functions, which is crucial for meeting regulatory capital requirements under frameworks such as Basel III. The Violations column for ES is marked as N/A because Expected Shortfall measures the magnitude of tail losses rather than the frequency of exceedances. Overall, the

acceptance of the GPD model underscores the robustness of EVT in capturing fat-tailed behaviour. These results align with advances in joint VaR–ES modelling [60, 61], reinforcing the GPD as a sophisticated tool for financial decision-making [62].

Economically, these results demonstrate a clear hierarchy of market soundness. The extreme tail thickness found in BTC/USD confirms that it functions primarily as a speculative vehicle; its "unbounded" downside risk poses a significant threat to capital preservation during market stress. In contrast, the BWP/USD exhibits the highest degree of market soundness; its thinner tails and near-zero shape parameters suggest a highly effective exchange-rate regime that successfully dampens extreme shocks. The ZAR/USD occupies a middle ground, showing moderate resilience but remaining susceptible to tail-risk spillovers from global commodity shifts. Collectively, these findings imply that while Bitcoin has gained institutional presence, the asymmetry between its gains and losses necessitates much higher prudential capital buffers than those required for traditional fiat currencies. For market stability, institutional portfolios and central bank reserves must be calibrated specifically to these asymmetric 99% Expected Shortfall levels to ensure they can absorb sudden, non-linear depreciations identified in the data.

5. Conclusion

This study provides empirical evidence supporting the superiority of the Generalised Pareto Distribution over the Generalised Extreme Value distribution when modelling extreme exchange rates, based on tail behaviour, for BTC/USD, ZAR/USD, and BWP/USD. Weekly extreme data reveal differing behaviours among each asset's tails. Although the GEV model produced positive shape parameters for all three assets, placing them in the Fréchet category of heavy-tailed distributions, the GPD offered a superior statistical fit. Especially for both ZAR/USD and BTC/USD, the GEV model is rejected at the 5% significance level in the lower tails; however, the GPD effectively modelled the unlimited nature of extreme outcomes observed in both assets.

These findings have important implications for macroeconomic stability in Southern Africa. Specifically, the presence of heavy-tailed risk in the ZAR/USD exchange rate, demonstrated by a negative shape parameter for tail values of $\xi = -0.2698$, suggests an increased probability of "black swan" type event occurring in response to either global commodity price shock or changes in sentiment toward emerging economies. As such, the SARB will be required to hold larger-than-normal foreign exchange reserves and develop new stress-testing procedures that account for risks associated with heavy-tailed devaluation scenarios. Conversely, while both the GPD estimated high shape parameters of $\xi_{loss} = 0.6494$ and $\xi_{gain} = 0.6697$ for both loss and gain cases, respectively, for BWP/USD, these symmetric and consistent estimates suggest that Botswana has maintained control over their managed exchange rate mechanism, with the Pula being pegged to a basket of currencies. These stable conditions support Botswana as a low-risk destination for foreign direct investment (FDI) and any identified tail thickness would most likely result from regional contagion from the more volatile South African economy versus internal structural instability. Finally, the analysis confirmed a large degree of asymmetry in the tail behaviour between downside and upside risks. The estimated lower-tail shape parameter value for Bitcoin of $\xi = -0.1138$ is approximately twice as large as the corresponding upper-tail shape parameter value of $\xi = 0.0612$, suggesting that downward crashes are much heavier-tailed than upward gains. Consistent with this observation, the risk metrics show that the downside VaR at a 99% confidence interval of 0.1888 is significantly greater than the expected upside potential, supporting the notion that Bitcoin is primarily used as a speculative tool. Therefore, regulatory bodies should require banks to hold additional capital to absorb the expected ES value of 0.2597.

Although this research presents a mathematical framework to evaluate risk, it suffers from two limitations: 1) the assumption of constant tail parameters; and 2) no consideration is made for liquidity risk. Given that financial markets can evolve quickly through dynamic volatility regimes, fixed estimates of location ($\hat{\mu}$), scale ($\hat{\sigma}$), and shape ($\hat{\xi}$) may not provide sufficient representation of time-varying characteristics of tail behaviour. Additionally, although this analysis used weekly extremes to evaluate long-term patterns of extreme movement, it does not explicitly quantify liquidity effects during periods of high stress. Research is now underway to evaluate deep learning models that incorporate elements of EVT. Current efforts examine whether hybrid/non-hybrid models can represent time-varying dynamics without relying upon assumptions inherent in classical parametric representations.

Acknowledgement

The authors gratefully acknowledge the constructive comments and suggestions provided by the anonymous reviewers, which have substantially improved the quality and clarity of this manuscript.

REFERENCES

1. Abubakari, M. (2023). Modelling exchange rate volatility in emerging markets: Evidence from the Ghana Cedi. *Journal of International Finance and Economics*, 23(1), 45–60.
2. Gyamerah, S. A. (2020). Modelling and forecasting the volatility of exchange rate using GARCH models. *International Journal of Business and Economics Research*, 9(1), 21–32.
3. Dinardi, F. B. (2019). Forecasting the Stock Market Using ARIMA and ARCH/GARCH Approaches. *Master's Dissertation*, Universidade NOVA de Lisboa, Lisbon, Portugal.
4. Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
5. Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
6. Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
7. Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779–1801.
8. Zakoian, J. M. (1994). Threshold heteroskedasticity models. *Journal of Economic Dynamics and Control*, 18(5), 931–955.
9. Klaassen, F. (2002). Improving GARCH volatility forecasts with regime-switching GARCH. *Empirical Economics*, 27(2), 363–394.
10. Enow, S. T. (2025). An empirical investigation of the long memory effect on the relationship between downside risk and stock returns: Evidence from international stock markets. *International Journal of Research in Business and Social Science*, 14(8), 213–219.
11. Lee, M. C. (2025). A Hybrid EGARCH–Informer Model with Consistent Risk Calibration for Volatility and CVaR Forecasting. *Mathematics*, 13(19), 3105.
12. Barz, T., and Nastansky, A. (2024). Challenges of financial risk management: An empirical study of the Value at Risk approach in stressful situations. *Working Paper No. 57*, Universität Potsdam.
13. Enow, S. T. (2023). Exploring Volatility Clustering in Financial Markets and Its Implications. *Journal of Economic and Social Development*, 10(2), 1–5.
14. Mínguez, R. (2025). Extreme Value Theory applications in finance and insurance. In *Advances in Quantitative Analysis of Finance and Accounting*, Vol. 15. Springer.
15. Metwane, M. K., and Maposa, D. (2023). Extreme Value Theory Modelling of the Behaviour of Johannesburg Stock Exchange Financial Market Data. *International Journal of Financial Studies*, 11(4), 130.
16. Ndlovu, T., and Chikobvu, D. (2023). Extreme Risk Comparison of Bitcoin/USD and ZAR/USD. *Risks*, 11(6), 100.
17. Makatjane, K. (2024). Statistical efficiency in block maxima extraction: A comparative study of emerging market return series. *Journal of Statistical Computation*, 12(2), 88–104.
18. Mosanawe, L., and Makatjane, K. (2025). Dynamic Volatility and Tail Risk in BTC, BWP, and ZAR Exchange Rates: Bayesian SARMA-GARCH with Skewed Error Distributions. *Statistics, Optimization and Information Computing*, 15(4), 2339–2366.
19. R Core Team (2024). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
20. Posit Team (2023). *RStudio: Integrated Development Environment for R*. Posit Software, PBC, Boston, MA. <http://www.posit.co/>
21. Coles, S. (2001). *An Introduction to Statistical Modelling of Extreme Values*. Springer, London. DOI: <https://doi.org/10.1007/978-1-4471-3675-0>
22. Stephenson, A. G. (2002). evd: Extreme Value Distributions. *R News*, 2(2):0.
23. Pfaff, B., Zivot, E., and McNeil, A. (2016). *evir: Extreme Values in R*. R package version 1.7-4.
24. Gilleland, E. and Katz, R. W. (2016). extRemes 2.0: An extreme value analysis package in R. *Journal of Statistical Software*, 72(8):1–39. DOI: <https://doi.org/10.18637/jss.v072.i08>
25. Chikobvu, D. and Ndlovu, T. (2023). The generalised extreme value distribution approach to comparing the riskiness of Bitcoin/US Dollar and South African Rand/US Dollar returns. *Journal of Risk and Financial Management*, 16(4):253. DOI: <https://doi.org/10.3390/jrfm16040253>
26. Ndlovu, T. and Chikobvu, D. (2024). Estimating extreme value at risk using Bayesian Markov regime switching GARCH-EVT family models. In *Cryptocurrencies - Financial Technologies of the Future*. IntechOpen. DOI: <https://doi.org/10.5772/intechopen.113042>
27. Fisher, R. A. and Tippett, L. H. C. (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(2):180–190. DOI: <https://doi.org/10.1017/S0305004100015681>
28. Pickands III, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics*, 3(1):119–131. DOI: <https://doi.org/10.1214/aos/1176343003>
29. Masilo, B. M. and Makatjane, K. (2025). Electricity demand forecasting of value-at-risk and expected shortfall: The South African context. *International Journal of Energy Economics and Policy*, 15(1):481–489. DOI: <https://doi.org/10.32479/ijeep.17056>
30. Özari, Ç., Eren, Ö., and Saygin, H. (2019). A new methodology for the block maxima approach in selecting the optimal block size. *Tehnički Vjesnik*, 26(5):1292–1296. DOI: <https://doi.org/10.17559/TV-20180529125449>

31. Araveeporn, A. (2025). Improved probability-weighted moments and two-stage order statistics methods of generalised extreme value distribution. *Mathematics*, 13(14):2295. DOI: <https://doi.org/10.3390/math13142295>
32. Rafiu, A. (2024). Statistical analysis of extreme events: A review of distributions and applications. *Journal of Applied Probability and Statistics*, 19(2):45–62.
33. Abdulali, B. A. A., Abu Bakar, M. A., Ibrahim, K., and Mohd Ariff, N. (2022). Extreme value distributions: An overview of estimation and simulation. *Journal of Probability and Statistics*, 2022(1):5449751. DOI: <https://doi.org/10.1155/2022/5449751>
34. Singirankabo, E. and Iyamuremye, E. (2022). Modelling extreme rainfall events in Kigali city using generalised Pareto distribution. *Meteorological Applications*, 29(4):e2076. DOI: <https://doi.org/10.1002/met.2076>
35. Chukwudum, Q. C., Mwitwa, P., and Mung'atu, J. K. (2020). Optimal threshold determination based on the mean excess plot. *Communications in Statistics-Theory and Methods*, 49(24):5948–5963. DOI: <https://doi.org/10.1080/03610926.2019.1624765>
36. Martins, A. L. A., Liska, G. R., Beijo, L. A., Menezes, f. S. d., and Cirillo, M. Â. (2020). Generalised Pareto distribution applied to the analysis of maximum rainfall events in Uruguiana, RS, Brazil. *SN Applied Sciences*, 2(9):1479. DOI: <https://doi.org/10.1007/s42452-020-03225-3>
37. Liu, B. and Ananda, M. M. (2024). The application of accumulation tests to peaks-over-threshold modelling. *Model Assisted Statistics and Applications*, 19(1):38–51.
38. Nieto, M. R., and Ruiz, E. (2016). Frontiers in VaR forecasting and backtesting. *International Journal of Forecasting*, 32(2), 475–501. <https://doi.org/10.1016/j.ijforecast.2015.06.005>
39. Engle, R. F., and Manganelli, S. (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4), 367–381. <https://doi.org/10.1198/073500104000000136>
40. Kupiec, P. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2), 73–84. <https://doi.org/10.3905/jod.1995.407942>
41. Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, 841–862. <https://doi.org/10.2307/2527344>
42. Papastathopoulos, I., and Tawn, J. A. (2013). Extended generalised Pareto models for tail estimation. *Journal of Statistical Planning and Inference*, 143(3), 131–143. <https://doi.org/10.1016/j.jspi.2012.07.006>
43. Fissler, T., and Ziegel, J. F. (2016). Higher order elicibility and applications to risk measures. *The Annals of Statistics*, 44(4), 1680–1707. <https://doi.org/10.1214/15-AOS1381>
44. Dimitriadis, T., and Bayer, S. (2019). A joint quantile and expected shortfall regression framework. *Journal of Financial Econometrics*, 171(2), 211–239. <https://doi.org/10.1093/jjfinec/nby004>
45. Chikobvu, D. and Jakata, O. (2020). Analysing extreme risk in the South African financial index (J580) using the generalised extreme value distribution. *Statistics, Optimization & Information Computing*, 8(4):915–933. DOI: <https://doi.org/10.19044/soic.v8i4.915>
46. Klarman, S. (2026). *Bitcoin's Volatility and Long-Term Investment Realities: Annualised Volatility Analysis*. ResearchGate Asset Allocation Insights.
47. Hassan, M. and Khan, A. (2026). Volatility in focus: comparing cryptocurrencies, fiat currencies from high-inflation economies, and gold. *Studies in Economics and Finance*, 43(2):354–372.
48. Arias, J. (2026). *Bitcoin's Role within an Institutional Portfolio: Data through February 2026*. Global Digital Asset Research Report.
49. Kayani, G. M. and Hasan, M. (2025). Cryptocurrency in portfolio management: Risk-return optimisation and diversification efficiency. *Journal of Economics and Finance*, 3(2):103–119.
50. Zhao, L., Wang, Q., and Zhang, H. (2025). Volatility Modelling and Tail-Risk Measurement in Cryptocurrency Markets: A GARCH-EVT approach. *Advances in Economics, Management and Political Sciences*, 224(1):92–98.
51. Bantsi, K. and Makatjane, K. (2026). Extreme energy prices forecasting: Application of alpha-recurrent neural network with generalised Pareto distribution. *South African Journal of Science*, 122(3/4): Art. #22326. <https://doi.org/10.17159/sajs.2026/22326>
52. Chan, K. and Lee, J. (2026). Stablecoin Shocks. *International Monetary Fund Working Paper*, WP/26/44.
53. Al-Maamary, F. (2025). Return and volatility spillover between cryptocurrencies, oil price and stock market in GCC countries. *Cogent Economics & Finance*, 13(1):2453584.
54. Naradh, K., Chinhamu, K., and Chifurira, R. (2021). Estimating the value-at-risk of JSE indices and South African exchange rate with Generalized Pareto and stable distributions. *Investment Management and Financial Innovations*, 18(3):151–165.
55. Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
56. Makatjane, K. and Shoko, C. (2025). Explainable Deep Learning for Financial Risk: Joint VaR and ES Forecasting Using ESRNN in the Bitcoin Market. *The African Finance Journal*, 27(1):53–69.
57. Schwaab, B. and Lucas, A. (2026). Tail risk and systemic linkages in global currency markets. *Journal of International Money and Finance*, 142:103045.
58. Melina, G. and Verona, F. (2024). Calibrating macro-prudential policy: A tail-risk approach. *European Economic Review*, 161:104612.
59. Fuentes, R. and Morales, M. (2025). Adapting to the extreme: Time-varying tail dynamics in cryptocurrency markets. *Quantitative Finance*, 25(3):412–428.
60. Hue, S. Y. and Biatour, B. (2024). Joint backtesting of VaR and Expected Shortfall: A scoring function approach. *Journal of Banking & Finance*, 158:107021.
61. Fu, Z. and Wang, Y. (2026). Advances in joint quantile and expected shortfall regression for financial risk management. *Computational Statistics & Data Analysis*, 191:107884.
62. Benito, S., López, J., and Romera, R. (2023). Assessing the importance of the choice threshold in quantifying market risk under the POT approach (EVT). *Risk Management*, 25(1):6. DOI: <https://doi.org/10.1057/s41283-022-00109-x>