



Dynamic Volatility and Tail Risk in BTC, BWP, and ZAR Exchange Rates: Bayesian SARMA-GARCH with Skewed Error Distributions

Letlhogonolo Mosanawe, Katleho Makatjane*

Department of Statistics, University of Botswana, Botswana

Abstract Exchange rate volatility presents significant risks to investors and governments, especially in developing economies and the cryptocurrency market, where unforeseen shocks may lead to considerable financial losses. Standard risk metrics frequently do not account for time-varying volatility, skewness, and fat-tailed return distributions, thereby constraining their predictive reliability. This research utilises a Bayesian SARMA–GARCH methodology with time-varying parameters to evaluate exchange rate risk for BTC/USD, BWP/USD, and ZAR/USD. Daily log returns are modelled with Asymmetric Generalised Error Distributions to address heavy-tailed and skewed characteristics. One-step-ahead forecasts of Value-at-Risk (VaR) and Expected Shortfall (ES) are produced and systematically backtested employing credible intervals, weighted continuous ranking probability scores, and dynamic quantile tests. The findings demonstrate substantial predictive accuracy, with Mean Prediction Interval Widths of 0.0518 for BTC/USD, 0.0722 for BWP/USD, and 0.0413 for ZAR/USD, and the majority of observed returns remaining within the 99% prediction intervals. BTC/USD responds rapidly to disturbances, ZAR/USD demonstrates persistent volatility, and BWP/USD reflects extended effects. The integration of time-varying dynamics and heavy-tailed distributions enhances the reliability of Value at Risk (VaR) and Expected Shortfall (ES) forecasts, thereby facilitating improved risk management, portfolio allocation, and regulatory oversight.

Keywords Cryptocurrency, Exchange rate, Foreign exchange markets, Heavy tails, Risk, Volatility clustering

AMS 2010 subject classifications 62-07, 62M10

DOI: 10.19139/soic-2310-5070-2883

1. Introduction

Cryptocurrencies are decentralised digital currencies that function independently of central banks or financial intermediaries by leveraging blockchain technology to enable secure transactions. Bitcoin, the leading cryptocurrency in terms of trading volume, possesses a market capitalisation of USD 452.1 billion. Although the adoption of Bitcoin may occur gradually, studies indicate that its popularity is expected to increase over time. According to [26], more than 2,300 US enterprises used Bitcoin and other digital assets in 2020 for diverse investment, operational, and transactional purposes [20]. Some economists contend that Bitcoin does not entirely satisfy the three conventional criteria of a currency: store of value, medium of exchange, and unit of account. However, it is exchangeable with fiat currencies and offers a more rapid method of transferring money [43]. Additionally, [42] noted that Bitcoin could improve lives and lower barriers to widespread participation, particularly if it were recognised as a legitimate currency.

Volatility represents an intrinsic attribute of financial assets and is a vital factor in risk management. Changes in returns and persistent volatility over time necessitate accurate modelling, as noted by [50]. Seminal models, including the autoregressive conditional heteroscedasticity (ARCH) developed by [28] and the generalised ARCH (GARCH) introduced by [15], were designed to capture these dynamics and are extensively utilised in forecasting

*Correspondence to: Katleho Makatjane (Email: makatjanek@ub.ac.bw). Department of Statistics, University of Botswana.

volatility for options pricing and risk management. Numerous GARCH specifications, as noted by [1, 37], utilise squared returns as the main determinant of conditional variance, as they yield an unbiased estimate of current volatility; however, this approach offers only a restricted indication of the actual level of market volatility. Financial systems are vulnerable to shocks and changing investor expectations. Research indicates that GARCH-family models with time-invariant parameters have limitations, as risk measures derived during stable periods may be biased in times of turbulence [47, 11]. Financial returns exhibit fat-tailed characteristics, necessitating a focus on tail risk, which poses measurement challenges under conventional parametric frameworks [48]. The advent of high-frequency data has facilitated the development of more precise volatility metrics, including realised volatility and realised kernels; however, these metrics are still susceptible to microstructure noise and non-trading hours [17, 24, 10]

Building on these insights, this study aims to forecast value-at-risk and expected shortfall up to 60 days ahead using daily currency data for Bitcoin/US Dollar, South African Rand/US Dollar, and Botswana Pula/US Dollar-adjusted closing prices. We employ a seasonal autoregressive integrated moving average (SARIMA) model with GARCH (SARIMA-GARCH) errors for out-of-sample prediction of VaR and ES. The SARIMA-GARCH framework allows the volatility parameters to evolve, capturing changing market conditions, structural breaks, and volatility clustering more accurately than conventional GARCH models. This approach provides reliable risk measures during periods of financial turmoil, with short-term forecasting having a superior impact on financial network stability [63]. A key factor in accurate SARIMA modelling is ensuring independently and identically distributed (i.i.d.) residuals. While previous studies have applied alternative methods, such as core vector regression with Gaussian process regression for short-term forecasting under uncertainty [19], or the price of volatile stocks (PVSt) as a risk proxy [57], our approach employs a conservative bootstrap method for probability forecasting as seen in the work of [14]. This extends conventional GARCH bootstrap estimates to SARIMA-GARCH VaR and ES calculations, capturing volatility clustering uncertainty more precisely. The rest of the paper is as follows: section 2 describes the methodology used in detail, section 3 presents the analysis and results, and section 4 discusses the findings of the paper and concludes.

Our methodology makes a notable contribution to VaR and ES forecasting. Building on [45], it incorporates real-time forecasting to detect sudden market shifts, particularly in highly dynamic economies. Unlike frequentist MLE approaches, we employ Bayesian estimation, which combines prior knowledge with observed data, is well-suited for small samples, and produces reliable forecasts. The procedure integrates threshold and quantile-weighted scoring methods, enabling rigorous comparison of VaR and ES density forecasts while providing robust tools for evaluating predictive performance. Additional innovations include the use of fixed-design residual bootstrap, tail risk (TR), and dynamic quantile (DQ) backtesting. The real-time forecasting system has been implemented on AWS, utilising EC2 instances for scalable computation and Lambda functions for serverless, event-driven predictions, ensuring rapid, reliable, and efficient deployment of risk forecasts.

Research Highlights and Key Findings

This study is one of the first to combine Bayesian SARMA-GARCH models with skewed error distributions to analyse time-varying volatility and tail risk in BTC/USD, BWP/USD, and ZAR/USD returns. By capturing downside risk in these currency markets, the study provides insights into extreme market movements beyond average returns. The highlights and key findings of this study are summarised in Table 1.

The rest of the paper is as follows: section 2 describes the methodology used in detail, section 3 presents the analysis and results, and section 4 discusses the findings of the paper and draws conclusions

Table 1. Performance and Contributions of the Bayesian SARMA-GARCH Model

Characteristic	Highlights	Findings	Contributions
Tail Risk Detection	Captures extreme market events using skewed GED distributions	BTC/USD reacts rapidly to shocks, ZAR/USD shows persistent volatility, BWP/USD shows prolonged effects	Identifies asymmetric and heavy-tailed behaviour in cryptocurrency and emerging FX markets
Error Distribution Analysis	Skewed GED outperforms symmetric GED and Student-t distributions	Accurately models both skewness and leptokurtosis	Improves prediction of extreme returns and enhances risk measurement precision
VaR and ES Back-testing	One-step-ahead density forecasts with bootstrapped uncertainty intervals	Forecasts are accurate, unbiased, and statistically reliable	Provides robust tools for regulatory compliance and portfolio risk management
Time-varying Volatility	Incorporates dynamic parameters in SARMA-GARCH framework	Effectively captures volatility clustering and evolving market conditions	Supports more responsive and adaptive risk management strategies
Methodological Innovation	Bayesian estimation with MCMC and bootstrapped credible intervals	Produces reliable parameter estimates and uncertainty ranges	Integrates probabilistic forecasting with extreme event modelling for improved decision-making

2. Methodology

Let r_t be stock returns at time t , [12] showed that r_t can be modeled by

$$r_t = \mu_t + \varepsilon_t \quad (1)$$

where μ_t is the time-varying mean; ε_t is the error term that can be modeled by

$$\varepsilon_t = \nu_t \sigma_t. \quad (2)$$

In model (2), σ_t is the time-varying dynamic, while ν_t is the i.i.d residuals from model (1).

2.1. Volatility Forecasting Models

In conventional SARIMA models, the disturbance term's variance is thought to be constant. The assumption of homoscedasticity (constant variance) seems inappropriate given that the data exhibit non-constant mean and variance as well as numerous seasons corresponding to weekly and monthly periodicity. The GARCH modelling technique is introduced to take the possibility of serial correlation in volatility into consideration. Over the past thirty years, a great deal of research has been done on the use of GARCH-type models to model volatility in time series data. Our work is closely connected to that of [63], who examined approaches for forecasting Net Imbalance Volume density. The problem was broken down by these writers into two parts: point forecasting and volatility forecasting. Good results were obtained using a seasonal ARMA model and a periodic AR model with simple volatility predictions.

Generalised Autoregressive Conditional Heteroscedasticity is a statistical model for evaluating time-series data in which the variance error is assumed to be serially autocorrelated. The study will employ the GARCH model to predict projected portfolio losses for risk management [59]. The GARCH models could handle volatility clustering over extended periods, resulting in a very strong sample of estimates. The [13] model, established by [15] as an extension of the ARCH model, estimates asset volatility. The parameters of the GARCH model are usually estimated using historical returns of an asset price, by making use of the maximum likelihood method as in the work of [66] and Bayesian approaches, just as in the work of [62]. In light of this, this study presents the combination of SARIMA(p,d,q)(P,D, Q) and GARCH(p,q) by referencing the research of [63], in which the authors forecasted

South Africa's daily peak power consumption. Our major goal, in this case, is to estimate the conditional mean of the returns process using the features of the SARIMA model and to estimate the conditional variance of the process using seasonalities and GARCH errors.

2.1.1. SARIMA–GARCH Model The SARIMA–GARCH model is one in which the variance of the error term from the SARIMA model follows a GARCH process. Specifically, the SARIMA component captures both non-seasonal and seasonal dynamics of the series, while the GARCH component accounts for time-varying volatility in the residuals. Following [67], the model can be expressed as follows

$$\begin{aligned} \Phi_P(B^s) \phi_p(B) (1 - B)^d (1 - B^s)^D X_t &= c + \Theta_Q(B^s) \theta_q(B) \varepsilon_t, \\ \varepsilon_t &= \sigma_t \nu_t, \quad \nu_t \sim \text{i.i.d.}, \quad E[\nu_t] = 0, \quad \text{Var}[\nu_t] = 1, \\ \sigma_t^2 &= \alpha_0 + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2. \end{aligned} \quad (3)$$

Here, X_t denotes the observed time series, $\phi_p(B)$ and $\theta_q(B)$ are the non-seasonal AR and MA polynomials of orders p and q , while $\Phi_P(B^s)$ and $\Theta_Q(B^s)$ are the seasonal AR and MA polynomials of orders P and Q with seasonality s . The terms d and D denote the non-seasonal and seasonal differencing orders, respectively. The residual ε_t represents the deviation of the series from its conditional mean, and σ_t^2 is the conditional variance of ε_t determined by the GARCH process.

2.2. Model Selection

The study uses the Akaike Information Criterion (AIC), Hannan–Quinn Information Criterion (HQIC), and the Bayesian Information Criterion (BIC) for lag length selection of the ADF model. According to [56], AIC is one of the most widely used and powerful criteria which is defined as follows

$$\text{AIC} = n \log \left(\frac{\text{RSS}}{n} \right) + 2k, \quad (4)$$

where n is the sample size, RSS is the residual sum of squares from the fitted ADF model, and k is the number of estimated parameters. The model with the lowest AIC is selected as the best-fitting model. Similarly, the BIC is given by

$$\text{BIC} = n \log \left(\frac{\text{RSS}}{n} \right) + k \log(n), \quad (5)$$

where n and k are defined as above. The model with the lowest BIC is considered optimal. The Hannan–Quinn criterion (HQ) is asymptotically efficient in estimating the best model because it provides fine-tuning even for large sample sizes, due to its $\log \log(n)$ term [38] shows that it is computed by

$$\text{HQ} = -2 \log \tilde{L}_n(\theta_n) + 2ck_0 \log \log n, \quad c > 1, \quad (6)$$

where k_0 is the number of estimated parameters, n is the sample size, and $\tilde{L}_n(\theta_n)$ is the likelihood function.

Following [16], the autocorrelation function (ACF), partial autocorrelation function (PACF), and cross-autocorrelation function (CACF) are used to tentatively identify SARIMA parameters p, q, P, Q . According [46], statistical measures provide additional evidence for selecting an appropriate intervention function; therefore theoretical behaviour of the ACF and PACF for a stationary SARMA process is summarised in Table 2. For robust results, the observed ACF and PACF patterns should align with these theoretical patterns [51].

[40] shows that the ACF can be derived as follows

$$\rho_k = \frac{\text{cov}(X_t, X_{t-k})}{\text{Var}(X_t)}, \quad k = 0, 1, 2, \dots \quad (7)$$

Table 2. Behaviour of Theoretical ACF and PACF for Stationary SARMA Processes

	SARMA(P,0)	SARMA(0,Q)	SARMA(P,Q)
ACF	Tails off at lags k	Cuts off after lag Q	Tails off at lags k
PACF	Cuts off after lag P	Tails off at lag k	Tails off at lag k

where the autocovariance is computed by

$$\text{cov}(X_t, X_{t-k}) = E(X_t^2) - (E(X_t))^2 = \text{Var}(X_t). \tag{8}$$

A correlogram is used to visualise ρ_k against k , hence the PACF is computed by using Model 9 as follows

$$a_{kk} = \begin{cases} \rho_1, & k = 1 \\ \frac{\rho_k - \sum_{j=1}^{k-1} a_{k-1,j} \rho_{k-j}}{1 - \sum_{j=1}^{k-1} a_{k-1,j} \rho_j}, & k = 2, 3, 4, \dots \end{cases} \tag{9}$$

2.2.1. *Bayesian SARIMA–GARCH* In the Bayesian framework, the likelihood of the SARIMA–GARCH model can be written in terms of the residuals ε_t and the conditional variance σ_t^2 . Let the residuals from the conditional mean be

$$\varepsilon_t = X_t - \mu_t, \tag{10}$$

where the conditional mean is

$$\mu_t = c + \sum_{i=1}^p \phi_i X_{t-i} + \sum_{j=1}^q \theta_j \varepsilon_{t-j} + \sum_{P=1}^P \Phi_P X_{t-sP} + \sum_{Q=1}^Q \Theta_Q \varepsilon_{t-sQ}. \tag{11}$$

The GARCH variance recursion is

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2, \tag{12}$$

which accounts for the persistence of shocks and time-varying volatility. To express the likelihood in a summation form, define the contributions from non-seasonal lags as

$$A = \sum_{t=1}^n \varepsilon_t^2, \quad B = \sum_{t=1}^n \sum_{i=1}^p \phi_i \varepsilon_t \varepsilon_{t-i}, \quad C = \sum_{t=1}^n \sum_{i=1}^p \sum_{j=1}^p \phi_i \phi_j \varepsilon_{t-i} \varepsilon_{t-j}, \tag{13}$$

and from seasonal lags as

$$S = \sum_{t=1}^n \sum_{P=1}^P \Phi_P \varepsilon_t X_{t-sP} + \sum_{t=1}^n \sum_{Q=1}^Q \Theta_Q \varepsilon_t \varepsilon_{t-sQ} \tag{14}$$

where the GARCH variance component is captured by

$$V = \sum_{t=1}^n \log(\sigma_t^2) + \sum_{t=1}^n \frac{\varepsilon_t^2}{\sigma_t^2}. \tag{15}$$

Thus, the SARIMA–GARCH likelihood, up to a proportionality constant, can be expressed as follows

$$L(\phi, \theta, \Phi, \Theta, \alpha, \beta \mid X_{1:n}) \propto \prod_{t=1}^n \sigma_t^{-1} \exp \left\{ -\frac{1}{2} \sum_{t=1}^n \frac{\varepsilon_t^2}{\sigma_t^2} \right\}. \quad (16)$$

Equivalently, the summation form mirroring classical AR(1) derivations is given by

$$\begin{aligned} \sum_{t=1}^n \frac{\varepsilon_t^2}{\sigma_t^2} = \sum_{t=1}^n \frac{1}{\sigma_t^2} & \left[(X_t - c)^2 - 2(X_t - c) \left(\sum_{i=1}^p \phi_i X_{t-i} + \sum_{P=1}^P \Phi_P X_{t-sP} \right) \right. \\ & \left. + \left(\sum_{i=1}^p \phi_i X_{t-i} + \sum_{P=1}^P \Phi_P X_{t-sP} \right)^2 + \left(\sum_{j=1}^q \theta_j \varepsilon_{t-j} + \sum_{Q=1}^Q \Theta_Q \varepsilon_{t-sQ} \right)^2 \right] \end{aligned} \quad (17)$$

and the log-likelihood becomes

$$\log L \propto -\frac{1}{2} \sum_{t=1}^n \left(\log(\sigma_t^2) + \frac{1}{\sigma_t^2} (A + B + C + S) \right). \quad (18)$$

Finally, the Bayesian posterior distribution is obtained by combining this likelihood with prior distributions on the parameters:

$$\pi(\phi, \theta, \Phi, \Theta, \alpha, \beta \mid X_{1:n}) \propto L(\phi, \theta, \Phi, \Theta, \alpha, \beta \mid X_{1:n}) \cdot \pi(\phi, \theta, \Phi, \Theta, \alpha, \beta), \quad (19)$$

where

$$\pi(\cdot)$$

are priors chosen to enforce stationarity, invertibility, and positivity of variance parameters. This formulation clearly separates the **SARIMA mean structure** from the **GARCH volatility structure**, allowing for full Bayesian inference using methods such as Markov Chain Monte Carlo.

2.3. Proposed Conditional Distributions

This section describes the four main univariate error distributions used in this study for modelling the markets' volatilities, and they include the GED, Skewed distributions by inverse scale Factors, generalised hyperbolic distribution, and the generalised hyperbolic skew student distribution.

2.3.1. Generalised Error Distribution The Generalised Error Distribution is a three parameter distribution belonging to the exponential family with conditional density given by

$$f(x) = \frac{\xi \exp 1/2 \left| \frac{x-\mu}{\sigma} \right|^\xi}{2^{1+\xi-1} \sigma \Gamma(\xi-1)}, \quad (20)$$

where, μ , ξ and σ are the location, shape and scale parameters respectively. The location parameter also corresponds to the distribution's mode, median, and mean since the distribution is symmetric and unimodal (i.e. μ). All odd moments beyond the mean are zero by symmetry. The variance and kurtosis is given by model (21) and Model (22) as

$$\text{Var}(x) = \sigma^2 2^{2/\xi} \frac{\Gamma(3\xi-1)}{\Gamma(\xi-1)} \quad (21)$$

and

$$\text{Kurt}(x) = \frac{\Gamma(5\xi-1) \Gamma(\xi-1)}{\Gamma(3\xi-1) \Gamma(3\xi-1)}. \quad (22)$$

When ξ declines, the density gets flatter and flatter until it approaches the uniform distribution as $\lim_{\xi \rightarrow \infty}$. When $\xi = 1$ and $\xi = 2$, the distribution becomes the Laplace and the normal, respectively. A unit standard deviation in Model (21) can be obtained by rescaling the density during standardisation, and this becomes

$$\text{Var}(x) = \sigma^2 2^{2/\xi} \frac{\Gamma(3\xi^{-1})}{\Gamma(\xi^{-1})} = 1. \quad (23)$$

When this is substituted into the scaled density of z , it becomes

$$f\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sigma} \frac{\xi \exp\left(-\frac{1}{2} \left| \sqrt{\frac{2^{-2/\xi} \Gamma(\xi^{-1})}{\Gamma(3\xi^{-1})}} \left(\frac{x-\mu}{\sigma}\right) \right|^{\xi}\right)}{\sqrt{2^{-2/\xi} \frac{\Gamma(\xi^{-1})}{\Gamma(3\xi^{-1})} 2^{1+\xi^{-1}} (\xi^{-1})}} \quad (24)$$

where, $\left(\frac{x-\mu}{\sigma}\right)$ is a standard normal distribution denoted Z .

2.3.2. Skewed Distributions by Inverse Scale Factors [31] suggested using inverse scale factors to add skewness to symmetric and unimodal distributions by adding them to the positive and negative real half lines. In light of a skew parameter, ξ , the density of a random variable Z can be represented as

$$f(x|\xi) = \frac{2}{\xi + \xi^{-1}} \left[f(\xi Z) H(-Z) + f(\xi^{-1} Z) H(Z) \right] \quad (25)$$

where, $\xi \in \mathfrak{R}$ and $H(\cdot)$ is the Heaviside function. The absolute moments required to derive the central moments are generated from the following function

$$M_r = 2 \int_0^{\infty} Z^r f(Z) dz. \quad (26)$$

The Normal, Student, and GED distributions have skew variants which have been standardised to zero mean and unit variance by using the moment conditions given by Model (27) and Model (28), respectively.

$$E[Z] = M_1 (\xi - \xi^{-1}) \quad (27)$$

and

$$\text{Var}[Z] = (M_2 - M_1^2) (\xi^2 + \xi^{-1}) + 2M_1^2 - M_1. \quad (28)$$

And for this study, a skewed Student's t-distribution, and skewed GED will be used.

Student's-t Distribution The Student's-t distribution introduced by [35] introduces a degrees-of-freedom parameter ν that governs tail thickness, making it suitable for modelling heavy-tailed financial returns. The conditional density of the residuals ε_t is given by:

$$f(\varepsilon_t | \nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\pi(\nu-2)} \Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{\varepsilon_t^2}{\nu-2}\right)^{-\frac{\nu+1}{2}}, \quad (29)$$

where $\Gamma(\cdot)$ is the gamma function. Smaller values of ν correspond to heavier tails, allowing the model to accommodate extreme returns that frequently occur in financial markets.

2.4. Parametric estimation of VaR and ES

To estimate VaR and ES using the parametric approach, two formulas can be used depending on the assumed distribution. A Gaussian distribution or a Student's t-distribution. Suppose the logarithmic returns, r , belong to a Gaussian distribution, an estimation of the next period, VaR and ES are expressed as

$$\widehat{VaR}_{t+1}^q = -\hat{\mu}_{t+1} + \hat{\sigma}_{t+1} \Phi^{-1}(1 - q) \quad (30)$$

$$\widehat{ES}_{t+1}^q = -\hat{\mu}_{t+1} + \hat{\sigma}_{t+1} \frac{\phi(\Phi^{-1}(1 - q))}{q} \quad (31)$$

where, $-\hat{\mu}_{t+1}$ is the expected loss in the coming period, $\hat{\sigma}_{t+1}$ is the returns' estimated standard deviation of the coming period, Φ^{-1} is the inverse cumulative distribution function of the standard Gaussian probability distribution $N(0, 1)$, and ϕ is the probability density function of the standard Gaussian probability distribution $N(0, 1)$.

Now, suppose the returns belong to a Student's t-distribution, The factor $\sqrt{\frac{v-2}{v}}$ is used to scale the distribution to give it a unit volatility. An estimation of the next period VaR and ES is expressed as

$$\widehat{VaR}_{t+1}^q = -\hat{\mu}_{t+1} + \hat{\sigma}_{t+1} \sqrt{\frac{v-2}{v}} (t_v^{-1}(1 - q)) \quad (32)$$

$$\widehat{ES}_{t+1}^q = -\hat{\mu}_{t+1} + \hat{\sigma}_{t+1} \sqrt{\frac{v-2}{v}} \cdot \frac{(g_v(t_v^{-1}(1 - q)))}{q} \cdot \frac{(v + (t_v^{-1}(1 - q)))^2}{v - 1} \quad (33)$$

where, $-\hat{\mu}_{t+1}$ and $\hat{\sigma}_{t+1}$ are the returns' estimated standard deviation of the coming period, v are the degrees of freedom and t_v^{-1} is the inverse cumulative distribution function of the Student's t-distribution and g_v is the probability density function of the Student's t-distribution with v degrees of freedom.

2.5. Fixed-Design Residual Bootstrap for Bayesian SARIMA–GARCH

Algorithm 1 describes a fixed-design residual bootstrap procedure used to estimate the distribution of conditional Value-at-Risk (VaR) and Expected Shortfall (ES) within a Bayesian SARIMA–GARCH framework. This algorithm incorporates both residual uncertainty and parameter uncertainty from the Bayesian posterior draws.

Algorithm 1: Fixed-Design Residual Bootstrap for Bayesian SARIMA–GARCH

for each posterior draw $\hat{\theta}_i$ **do**

for $t = 1, \dots, n$ **do**

 Generate bootstrap innovation as follows

$$\eta_t^* \sim \text{i.i.d. from } \hat{\mathbb{F}}_n, \quad (34)$$

 and construct the bootstrap residual by computing

$$\varepsilon_t^* = \tilde{\sigma}_t(\hat{\theta}_i) \eta_t^*. \quad (35)$$

end

Now, compute the bootstrap parameter estimate by using

$$\hat{\theta}_i^* = \arg \max_{\theta \in \Theta} L_n^*(\theta), \quad (36)$$

where the bootstrap log-likelihood is computed as follows

$$L_n^*(\theta) = \frac{1}{n} \sum_{t=1}^n \ell_t^*(\theta), \quad \ell_t^*(\theta) = -\frac{1}{2} \left(\frac{\varepsilon_t^*}{\tilde{\sigma}_t(\theta)} \right)^2 - \log \tilde{\sigma}_t(\theta). \quad (37)$$

for $t = 1, \dots, n$ **do**

 Compute the bootstrap standardised residual, which is given by

$$\hat{\eta}_t^* = \frac{\varepsilon_t^*}{\tilde{\sigma}_t(\hat{\theta}_i^*)}. \quad (38)$$

end

Finally, compute bootstrap VaR at level α as

$$\hat{\xi}_{n,\alpha}^* = \arg \min_{Z \in \mathbb{R}} \frac{1}{n} \sum_{t=1}^n \rho_\alpha(\hat{\eta}_t^* - Z), \quad (39)$$

where ρ_α is the check function used in quantile regression.
and Expected Shortfall as follows

$$ES_\alpha^* = \frac{1}{N_{\text{post}}} \sum_{i=1}^{N_{\text{post}}} \mathbb{E} \left[r_t(\ell) \mid r_t(\ell) < \hat{\xi}_{n,\alpha}^* \right]. \quad (40)$$

end

Remark

The term *fixed-design* refers to generating bootstrap residuals using the original conditional variances $\tilde{\sigma}_t(\hat{\theta}_i)$, which are functions of the observed residuals $\varepsilon_1, \dots, \varepsilon_{t-1}$ and the posterior parameter draw $\hat{\theta}_i$. In contrast, the recursive-design bootstrap recalculates the conditional variances dynamically from the bootstrap residuals themselves. The recursive design is more computationally intensive. For details, see Appendix B of [14]. Figure 1 gives a flow of Algorithm 1.

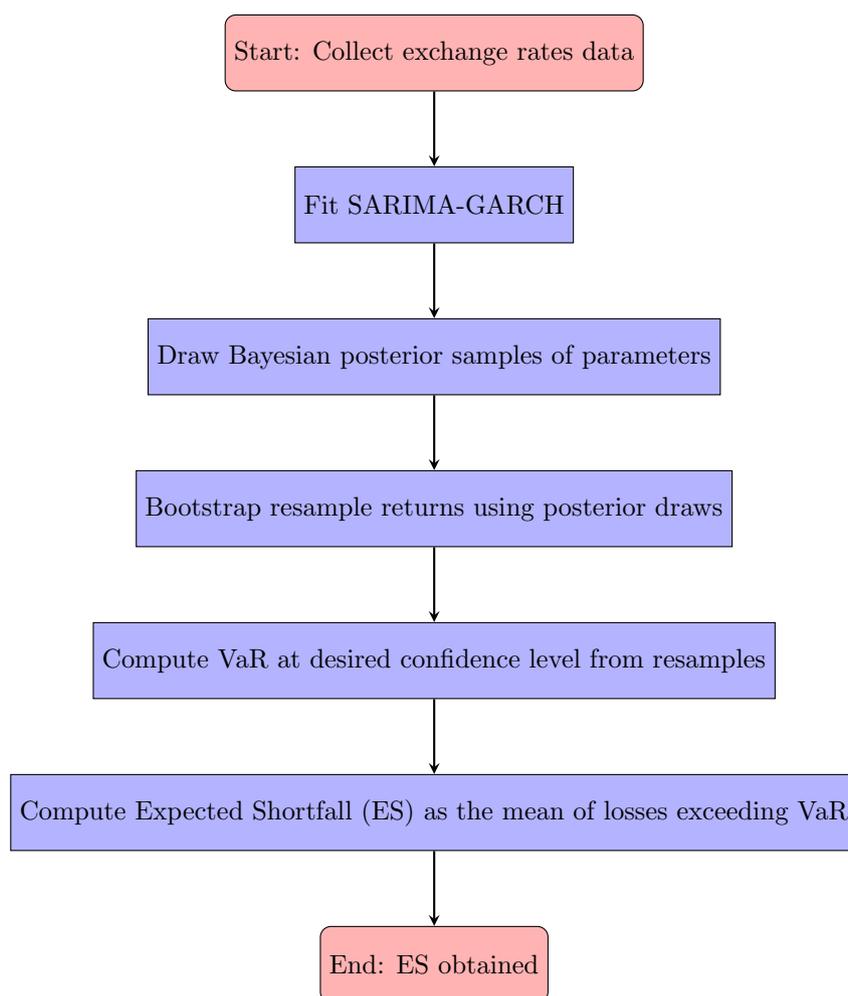


Figure 1. Bootstrap-based VaR and ES computation using Bayesian posterior parameter draws

2.6. Value-at-Risk and Expected Shortfall Forecasts Validation

The previous backtesting techniques simply consider the quantity of VaR and ES exceptions, ignoring their sizes. Conversely, [54] challenges the conditional coverage of [22] by arguing that the two-tailed nature of this test makes it capable of rejecting risk models that are too conservative. It is important to remember, nevertheless, that risk models may also be disapproved if they are excessively cautious, as this is undesirable for financial institutions. An alternative statistic is the tail risk (TR) statistic, which is described as

$$TR = -\frac{1}{T} \sum_{i=1}^T (R_t - \alpha) I(R_t - \alpha). \quad (41)$$

Model (41) is a TR statistic that provides risk managers with an estimate of the total tail losses that a portfolio could experience throughout the time under consideration. Assuming normal returns, the asymptotic distribution of the TR statistic is obtained.

The dynamic quantile test by [30] assesses the joint hypothesis that $E[d_t] = \alpha$ and the hit variables are independently distributed. The implementation of the test involves no mean process in $Hit_t^\alpha \equiv d_t - \alpha$. When a

model is correctly specified, $\bar{Hit}_t^\alpha = 0$ and $\text{Corr}(Hit_t^\alpha) = 0$. The DQ test is then the traditional Wald test of the joint nullity of all coefficients in the following linear regressions

$$Hit_t^\alpha = \begin{cases} \delta_0 + \sum_{t=1}^L \delta_1 Hit_{t-1}^\alpha + \delta_{L+1} VaR_{t-1}(\alpha) + \varepsilon_t \\ \delta_0 + \sum_{t=1}^L \delta_1 Hit_{t-1}^\alpha + \delta_{L+1} ES_{t-1}(\alpha) + \varepsilon_t. \end{cases} \tag{42}$$

With $L + 2$ degrees of freedom, model (42) has an asymptotic chi-square distribution under the null hypothesis of accurate unconditional and conditional coverage. The conventional choice is now $L = 4$ lags, as established by [30].

2.7. Proposed Threshold and Quantile-Weighted Scoring Rules

We examine density predictions in a time series context for the VaR and ES. A rolling window of the previous b observations is utilised to fit a density forecast for a future observation that is k time steps ahead. We consider $\omega_1, \dots, \omega_T$ as a stochastic process that can be further subdivided into $\omega_t = (X_t, r_t)$, where r_t is the variable of interest, which in this case is the returns series on Bitcoin/US dollar, Botswana Pula/US dollar, and South African Rand/US dollar. This is by the work of [45]. Let $T = m + n$; at time $t = m, \dots, m + n - k$, density forecasts \hat{f}_{t+k} and \hat{g}_{t+k} are generated, and they both rely on $\omega_{t-m+1}, \dots, \omega_t$. In this situation, the sole criterion for constructing a prediction is that it be a measurable function of the data in the estimation of a rolling window [7]. Nonetheless, [25] argued that rational users should select the optimal predictor, independent of the cost-loss structure. As a result, it is critical to establish the evaluation rule in the following sense

$$E_f S(f, Y) = \int f(y) S(f, d) dy, \quad \forall f, g. \tag{43}$$

If and only if $f = g$ and model (43) hold, then the scoring rule is strictly correct. Furthermore, a comparison of the average scores of the density forecast processes is made to rank them as follows

$$S_n^{-f} = \frac{1}{n - k + 1} \sum_{t=m}^{m+k-1} S(\hat{f}_{t+k}, y_{t+k}) \tag{44}$$

and

$$S_n^{-g} = \frac{1}{n - k + 1} \sum_{t=m}^{m+k-1} S(\hat{g}_{t+k}, y_{t+k}), \tag{45}$$

where f is preferred if $S_n^{-f} < S_n^{-g}$ and prefer g otherwise. [25] considered tests of equal forecast performance based on the test statistic in model (46) as

$$t_n = \sqrt{n} \frac{S_n^{-f} - S_n^{-g}}{\hat{\sigma}_n}, \tag{46}$$

where

$$\hat{\sigma}_n^2 = \frac{1}{n - k + 1} \sum_{j=-(k-1)}^{k-1} \sum_{t=m}^{m+n-k+|j|} \Delta_{t,k} \Delta_{t+|j|,k} \tag{47}$$

and $\Delta_{t,k} = S(\hat{f}_{t+k}, y_{t+k}) - S(\hat{g}_{t+k}, y_{t+k})$ as proposed by [4]. What scoring rule should be considered? To answer this question, we apply a weighted logarithmic scoring rule in model (48)

$$S(f, y) = w \left(\frac{y - \mu}{\sigma} \right) S_0(f, y) \tag{48}$$

to get a fixed, non-negative weight function, w where the estimations of the prediction of the unconditional mean and standard deviation are denoted by μ and σ , respectively. The previous m observations and the scoring criteria

represented by $S_0(f, y) = |\log f(y)|$ serve as the foundation for these non-negative functions. The weight function highlights important areas, such as a variable's centre of interest or tails. Denoting a standard normal probability density and cumulative distribution functions with ϕ and Φ , respectively, the weight functions $w_1(x) = \phi(x)$, $w_2(x) = 1 - \phi(x)\phi(0)$, $w_3 = \Phi(x)$, $w_4 = 1 - \Phi(x)$ emphasise the centre, the tails, the right tail, and the left tail

While the weighting strategy seems nice, it leads to incorrect conclusions and corresponds to the use of incorrect scoring criteria. In this study, the SARIMA–GARCH method is utilised rather than the GARCH(1,1) procedure, as in the work of [33]. Our objective is to develop a test that utilises the weighting strategy of [45]. By doing so, we want to prevent erroneous conclusions and provide corresponding graphical tools that can be utilised to identify the advantages and disadvantages of a forecasting technique. Although it is based on correctly weighted versions of a continuous ranked probability score (CRPS) [7], the test statistic from model (46) is kept. Any density forecast f induces a probability forecast for the binary event $\{X \leq Z\}$ passing through the value of the corresponding cumulative distribution function (CDF) in model (49) as

$$F(z) = \int_{-\infty}^z f(y)dy, \quad (49)$$

at the threshold of $Z \in \mathfrak{R}$. Similarly, this induces the quantile forecast $F^{-1}(\alpha)$ at the level $\alpha \in (0, 1)$. The continuous ranked probability score is now defined by

$$CRPS(f, y) = \int_{-\infty}^{\infty} PS(F(z), I\{y \leq z\}) dz \quad (50)$$

that can be simplified to

$$CRPS(f, y) = \int_0^1 QS_{\alpha}(F^{-1}(\alpha), y) d\alpha, \quad (51)$$

where $PS(F(z), I\{y \leq z\}) = (F(z) - I\{y \leq z\})^2$ is the Brier likelihood score for the probability forecast $F(z)$ of a binary event $\{Y \leq z\}$ at the threshold $z \in \mathfrak{R}$ and $QS_{\alpha}(F^{-1}(\alpha), y) = 2(I\{y \leq F^{-1}(\alpha)\} - \alpha)(F^{-1}(\alpha) - y)$ is the quantile score for the quantile forecast $F^{-1}(\alpha)$ at the level $\alpha \in (0, 1)$. Here, the symbol I stands for an indicator function. Following [33], a threshold-weighted version of the continuously ranked probability score is obtained as

$$S(f, y) = \int_{-\infty}^{\infty} PS(F(z), I\{y \leq z\}) u(z) dz, \quad (52)$$

where u is a non-negative weight function on the real line. In the same way, a quantile weighted version is obtained by

$$S(f, y) = \int_0^1 QS(F^{-1}(\alpha), y) v(\alpha) d\alpha \quad (53)$$

where a weight function on the unit interval, v , is non-negative. The suggested weight functions for the threshold and quantile weighted variants of the continuous ranking probability score are shown in Table 3. The probability density function $\phi_{a,b}$ and the cumulative distribution function $\Phi_{a,b}$ of the normal distribution with mean a and standard deviation b are used to specify the threshold weight functions. Table 3 shows Proposed Threshold and Quantile Weights. These weights are used to compare the density forecasts of both time-varying VaR and ES.

Table 3. Proposed Threshold and Quantile Weights

Threshold Weights	Emphasis	Quantile Weights
$\omega(Z) = \phi_{\{a,b\}}(Z)$	Center	$\omega(\alpha) = \alpha(1 - \alpha)$
$\omega(Z) = 1 - \phi_{\{a,b\}}(Z)$	Tails	$\omega(\alpha) = (2\alpha - 1)^2$
$W(Z) = \Phi_{\{a,b\}}(Z)$	Left Tail	$\omega(\alpha) = \alpha^2$
$\omega(Z) = 1 - \Phi_{\{a,b\}}(Z)$	Right Tail	$\omega(\alpha) = (1 - \alpha)^2$

3. Analysis and Results

This section covers an empirical analysis of the research using Currency data from January 2, 2015, to June 25, 2024. It was retrieved using the Python Yahoo Finance module and accessed on June 26, 2024. The adjusted closing values of daily exchange rates are fitted to the SARIMA-GARCH model with a student-t innovation. Bitcoin is traded every day, hence there are 3463. The rand and pula are not traded on weekends and South Africa’s and Botswana’s public holidays, resulting in 2469 observations for both BWP/US dollar and ZAR/US dollar. To align our data for analysis, we replaced missing values in the rand and pula exchange rates with zero since there are no profits or losses realised by the holder of the local currency during the weekend and public holidays as [20] as suggested. Plots of the adjusted closing stock prices are displayed in Figure2-Figure 4. Although the data is shown to be non-normal in Figure2(c and d), Figure 3(c and d), and Figure 4(c and d), the kernel density plot suggests that the distribution of the closing stock prices is leptokurtic. Notably, seasonality is associated with specific positive and negative patterns in Figure2(a), Figure 3(a), and Figure 4(a). Events like the COVID-19 epidemic and is to blame for these clusters of volatility. The largest concentrated return losses are also seen in these currencies. In addition, these Figures show a possible advantage when conditional heteroscedasticity is taken into account. This study emphasises two crucial elements: the cause of weight loss and the unpredictable nature of weight reduction. While the former argues that downturn volatility follows these shocks rather than large losses or gains, the latter argues that irregular shocks in the real business sector have a stronger effect on future volatility. There are isolated exceptional returns, which result from financial market shocks.

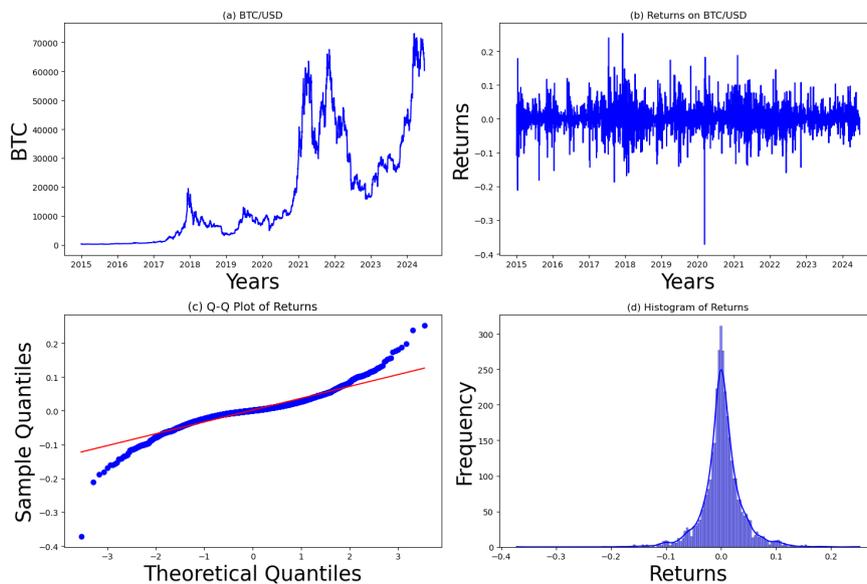


Figure 2. Plots of adjusted closing prices for BTC/USD

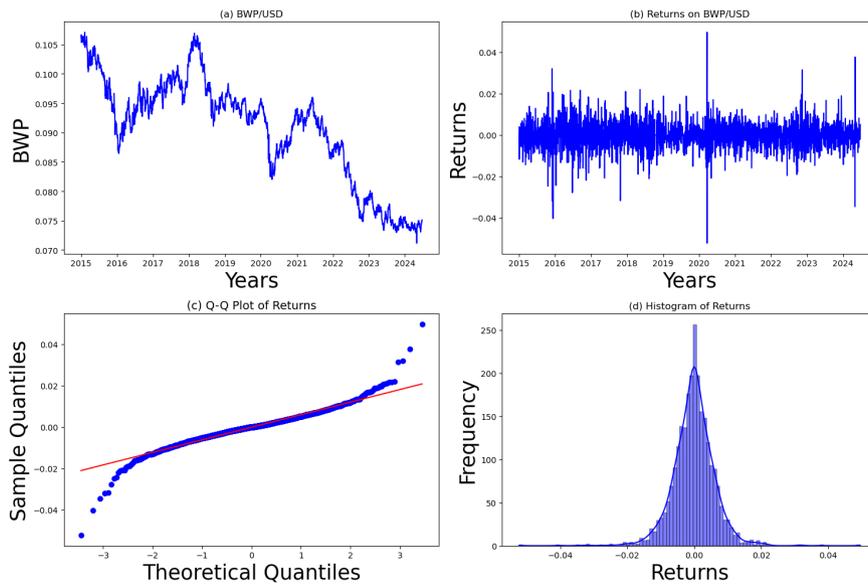


Figure 3. Plots of adjusted closing prices for BWP/USD

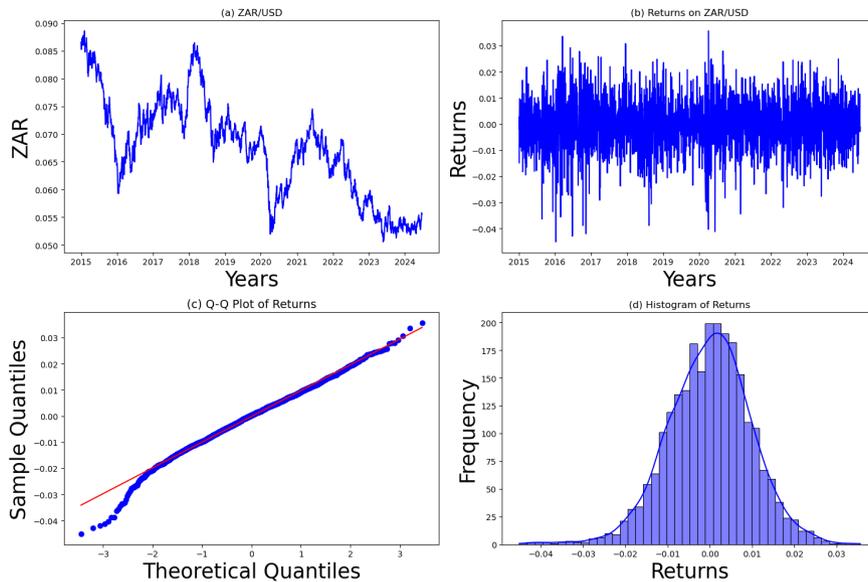


Figure 4. Plots of adjusted closing prices for ZAR/USD

The null hypothesis of normality using Jarque-Bera is rejected in Table 4 at the 5% level of significance, suggesting that symmetric models are not appropriate for evaluating the previously indicated return series. An insignificant p-value is obtained from the Ljung-Box test for ZAR/USD returns, suggesting that the null hypothesis of no autocorrelation is not rejected. This indicates that independent and identically distributed observations are anticipated (i.i.d). This null hypothesis is rejected for the BTC/USD and BWP/USD returns, though, and the autocorrelation problem will be resolved using the SARMA approach.

Table 4. Descriptive statistics of exchange rate price returns.

Currency	Observations	Mean	Median	Skewness	Kurtosis	Jarque-Bera
BTC/USD	3462	0.002199	0.00139	-0.1507	7.616	8378.633(0.0001)
ZAR/USD	2469	0.00013	0.00032	-0.29284	4.644	121.026(0.0001)
BWP/USD	2469	0.000121	0.00018	-0.1677	6.642	4548.543(0.0001)

NB: values in (.) are the p-values of the J-B Test

3.1. SARMA-GARCH Model Results

Parameter estimates are obtained using the `BayesGARCH` R package of [6], which builds on the methods of [27, 53]. We fit 12 competing SARIMA–GARCH models, ranging from SARIMA(1,0,1)(1,0,1)₇–GARCH(1,1) to SARIMA(12, 0, 12)(12, 0, 12)₇–GARCH(1,1), and select the optimal specification for each series using the Akaike Information Criterion (AIC) which was developed by [2], Bayesian Information Criterion (BIC) of [61], and Hannan–Quinn Criterion (HQC) of [36]. A 7-day seasonal period is adopted, as Bitcoin is traded daily. In contrast, for BWP/USD and ZAR/USD, weekends and public holidays are non-trading days. In line with the methodology of [20], these missing values are replaced with zeros, reflecting the absence of profits or losses on such days. Based on this model selection procedure, daily returns of BWP/USD and BTC/USD are characterised by a SARMA(0, 0, 0)(3, 0, 0)₇–GARCH(1,1) process, while ZAR/USD follows a SARMA(0, 0, 0)(4, 0, 0)₇–GARCH(1,1) process, because these models are the ones with the lowest AIC, BIC, and HQC values. The SARMA component is used to simulate the conditional mean dynamics of the return series, encapsulating short-term autocorrelation and seasonal patterns resulting from trading activity, market microstructure, and calendar impacts. The SARMA structure addresses predictable return patterns, whereas the GARCH component specifically simulates time-varying volatility, a common characteristic of financial returns. This separation guarantees that the residuals from the mean model resemble white noise, enabling the GARCH–skewed–GED component to concentrate on heteroskedasticity and extreme tail behaviour. Integrating SARMA for the mean and GARCH for the variance enhances the precision of Value-at-Risk and Expected Shortfall projections, making the model especially appropriate for portfolio risk management applications. The return series shows seasonal trends and short-term autocorrelation, indicating market activity. These results are reported in Figure 8, Figure 9, and Figure 10 in the Appendix section. Bitcoin trades every day, but BWP/USD and ZAR/USD do not on weekends or holidays. This causes systematic return changes that SARMA seasonal delays can capture. The ZAR/USD and BWP/USD show day-of-the-week influences, like developing and established currency markets. Seasonal autoregressive elements are included in the model because conditional mean dynamics mirror monthly and yearly impacts, albeit less so. The SARMA–GARCH specification explicitly models these patterns to produce volatility estimate residuals closer to white noise, boosting Value-at-Risk and Expected Shortfall projections.

Posterior distributions of the GARCH parameters are obtained via Markov Chain Monte Carlo (MCMC) sampling, using four independent chains of 50,000 iterations each, with the first 10,000 iterations discarded as burn-in. Prior distributions are specified as $\alpha \sim \mathcal{B}(2, 5)$, $\beta \sim \mathcal{B}(5, 1)$, and $\omega \sim \Gamma(2, 0.1)$, where the Beta priors for α and β restrict their values to (0, 1), while the Gamma prior for ω guarantees positivity. For BTC / USD, the SARMA coefficients in Table 5 indicate strong short-term dependence, with $\Phi_1 = -0.8737$, $\Phi_2 = -0.7075$, and $\Phi_3 = -0.5315$ with all $p < 0.001$, suggesting that past returns exert a significant influence on current returns.

The GARCH parameters have significant volatility clustering, evidenced by the sum $\alpha_1 + \beta \approx 0.8963$, which approaches unity. This indicates that volatility shocks are very enduring and need an extended period to fade. [34] in their study on an empirical study of volatility in the Cryptocurrency market also found the same results as found in this study. In practical terms, phases of elevated volatility are often followed by more volatility, whereas tranquil times tend to endure, illustrating the clustering tendency characteristic of financial time series. The minimal value of α_1 indicates that the initial response of volatility to new shocks or information is constrained, indicating that markets adapt incrementally to new data. Conversely, the substantial β value indicates that historical volatility has a significant and enduring impact on present volatility levels. This persistence may be attributed to reasons

Table 5. SARMA(0, 0, 0)(3, 0, 0)₇-GARCH(1,1) Coefficients for BTC/USD

	Estimate	Std. Error	t value	p-value
Φ_1	-0.8737	0.0167	-52.317	0.0001
Φ_2	-0.7075	0.0215	-32.907	0.0001
Φ_3	-0.5315	0.0229	-15.349	0.0001
α_0	0.1027	0.0098	10.479	0.0001
α_1	0.000015	0.000004	3.750	0.0008
β	0.8962	0.0102	87.862	0.0001
ν	1.0183	0.0204	49.916	0.0001

like speculative trading, herd behaviour among investors, and fluctuations in market moods, which lead to gradual volatility rather than abrupt reactions to individual occurrences [39].

Table 6. Error Distribution for BTC/USD

Distribution	MSE	RMSE	MAE	MFP
GED	0.000492	0.022192	0.020298	0.019783
Skewed GED	0.000153*	0.012391*	0.009210*	0.025732
Skew-student t	0.002251	0.047449	0.046384	-0.046372*
t-dist	6.47772	4.52523	4.52522	4.52522

The error distribution results in Table 6 indicate that the skewed Generalised Error Distribution (GED) provides the best fit, with the lowest MSE (0.000153), RMSE (0.012391), and MAE (0.009210). This suggests that BTC/USD returns are both asymmetric and heavy-tailed, characteristics that symmetric distributions fail to capture accurately. The standard GED and skewed Student-t distributions show higher errors, while the symmetric Student-t distribution performs poorly, significantly underestimating extreme movements. This highlights the importance of accounting for skewness and tail risk in modeling Bitcoin returns, which is particularly relevant for accurate VaR and ES forecasting.

Table 7. SARMA(0, 0, 0)(4, 0, 0)₇-GARCH(1,1) Coefficients for ZAR/USD

	Estimate	Std. Error	t value	p-value
Φ_1	-0.8123	0.0199	-40.819	0.0001
Φ_2	-0.6614	0.0250	-26.456	0.0001
Φ_3	-0.5177	0.0264	-19.609	0.0001
Φ_4	-0.6091	0.0250	-24.364	0.0001
α_0	0.0279	0.0012	23.250	0.0001
α_1	0.3186	0.237	1.344	0.0001
β	0.7999	0.0651	12.288	0.0001
ν	1.564	1.07	1.462	0.0001

The SARMA(0, 0, 0)(4, 0, 0)₇ coefficients for ZAR/USD presented in Table 7 indicate moderate persistence, with values of $\Phi_1 = -0.8123$, $\Phi_2 = -0.6614$, $\Phi_3 = -0.5177$, and $\Phi_4 = -0.3091$, each exhibiting p-values below 0.01%. The GARCH parameters' sum ($\alpha_1 + \beta \approx 1.1185$) indicates significant volatility persistence, suggesting that shocks to the ZAR/USD exchange rate have enduring and magnifying effects. When $\alpha_1 + \beta > 1$, the model suggests an explosive or non-stationary volatility process, indicating that volatility fails to revert to its long-term mean and instead accumulates over time [34]. This behaviour is frequently linked to markets experiencing structural changes, speculative trading, or prolonged uncertainty. The ZAR/USD's notable persistence likely indicates the Rand's responsiveness to global commodity price changes, variations in risk sentiment, and capital flow instability

in emerging markets. The dynamics reflect the influence of investor expectations and external shocks on the South African Rand. Evidence of near- or above-unity persistence has been observed in emerging markets and high-volatility assets, indicating non-mean-reverting or long-memory volatility structures [23, 9, 29].

With regards to the error distributions, a similar pattern is observed as reported in Table 8. The skewed GED again yields the lowest error metrics (MSE = 0.000113, RMSE = 0.002391, MAE = 0.002010), demonstrating its superior ability to capture both skewness and leptokurtosis in the return distribution. Other distributions, including the standard GED and skewed Student-t, show noticeably higher errors, while the symmetric Student-t is unsuitable due to its inability to model extreme returns accurately. These results suggest that ZAR/USD returns exhibit asymmetric heavy tails, and correctly modelling this feature is crucial for reliable risk measurement.

Table 8. Error Distribution for ZAR/USD

Distribution	MSE	RMSE	MAE	MFP
GED	0.004492	0.032192	0.070298	0.039783
Skewed GED	0.000113*	0.002391*	0.002010*	-0.005732*
Skew-student t	0.022251	0.047449	0.146384	0.016372
t-dist	5.47772	4.52523	4.52522	4.52522

The SARMA(0, 0, 0)(3, 0, 0)₇ coefficients demonstrate significant short-term dependence, with $\Phi_1 = -1.4109$, $\Phi_2 = -1.4120$, and $\Phi_3 = -1.1533$, which all have p -values < 0.001 , indicating that previous returns have a substantial impact on current returns. Volatility exhibits significant persistence, as indicated by the relationship $\alpha_1 + \beta \approx 1.1442$. The sum exceeding unity indicates a non-stationary or explosive volatility process in the model. This suggests that shocks to the BWP/USD exchange rate result in enduring and potentially amplifying effects, rather than a rapid return to a long-term mean. This notable persistence likely indicates an increased sensitivity to macroeconomic news, investor expectations, and external shocks impacting the Botswana Pula. Despite notable volatility, some research indicates that its direct effects on economic outcomes may be minimal. For instance, [60] demonstrates that while the Pula/Rand exchange rate shows volatility, it does not significantly influence Botswana's economic growth. Comparable patterns have been identified in additional emerging markets, where extreme volatility persistence is recorded; however, the influence on macroeconomic performance may be moderated by various factors [60]. These results show significant volatility in the BWP/USD exchange rate, primarily reflecting market dynamics and investor expectations rather than causing immediate economic consequences.

Table 9. SARMA(0, 0, 0)(3, 0, 0)₇-GARCH(1,1) Coefficients for BWP/USD

	Estimate	Std. Error	t value	p-value
Φ_1	-1.411	0.019	-74.263	0.00024
Φ_2	-1.412	0.031	-45.548	0.0001
Φ_3	-1.153	0.035	-32.943	0.0005
α_0	0.0005	0.015	0.033	0.0001
α_1	0.3476	0.262	1.327	0.0008
β	0.7966	0.088	9.052	0.0001
ν	1.517	0.152	9.980	0.0001

In the case of the Botswana Pula/US Dollar (BWP/USD) exchange rate, the skewed GED distribution again yields the most accurate forecasts, with MSE = 0.000053, RMSE = 0.000391, and MAE = 0.000210. This confirms that BWP/USD returns are highly skewed and heavy-tailed, making the skewed GED the most suitable choice for capturing extreme movements. Other distributions, particularly the symmetric Student-t, exhibit substantially higher errors, highlighting their inadequacy in modelling the true volatility structure of BWP/USD returns. These results are presented in Table 10

Table 10. Error Distribution for BWP/USD

Distribution	MSE	RMSE	MAE	MFP
GED	0.000492	0.022192	0.020298	0.019783
Skewed GED	0.000053*	0.000391*	0.000210*	-0.000732*
Skew-student t	0.005251	0.057447	0.046384	0.076372
t-dist	3.48772	4.52023	4.57889	4.6707

The observation that GARCH parameters reveal significant time-varying volatility in BTC/USD, ZAR/USD, and BWP/USD indicates that exchange rate returns do not adhere to a constant variance process. The rapid response of BTC/USD to shocks indicates a high sensitivity of the cryptocurrency market to new information or market events, highlighting its speculative characteristics. Such behaviour indicates that risk measures such as VaR and ES must incorporate rapid volatility fluctuations in Bitcoin trading, as neglecting these factors may lead to an underestimation of potential losses in short-term positions. The pronounced persistence in volatility for ZAR/USD and BWP/USD suggests that shocks to the exchange rate exhibit prolonged effects, a characteristic commonly observed among emerging-market investors. Policymakers should expect high volatility periods to persist, which necessitates increased vigilance during turbulent periods. Hedging strategies and capital allocation decisions must account for this extended risk exposure. The BWP/USD demonstrates significant volatility effects, indicating that shocks exert a persistent influence over time. This highlights the necessity for strong risk management structures, as conventional models that assume constant or rapidly diminishing volatility may fail to adequately account for tail risks in the Pula market. The results of the error distribution indicate that both cryptocurrencies and fiat currencies exhibit heavy-tailed and skewed return distributions, making standard symmetric distributions inadequate for modelling extreme events. The skewed GED distribution, which best aligns with the data, facilitates more precise predictions of infrequent yet significant market fluctuations. The result has direct implications for regulatory oversight, portfolio risk management, and derivative pricing, as traditional models may underestimate the probability and severity of extreme losses.

3.2. Testing stationarity and stability of the estimated SARIMA-GARCH-Skewed-GED

Finally, the stationarity analysis of the seasonal AR components indicates that BTC/USD exhibits roots with moduli greater than one, confirming a stationary seasonal structure. In contrast, ZAR/USD and BWP/USD have at least one root with modulus less than one, suggesting high persistence or potential non-stationarity in the seasonal components. This aligns with the pronounced volatility clustering observed in the GARCH estimates, where shocks to ZAR/USD and BWP/USD returns tend to persist for extended periods. The results imply that while BTC/USD returns are mean-reverting in the seasonal dimension, the emerging-market currencies exhibit more enduring seasonal effects, reflecting heightened sensitivity to macroeconomic news and market conditions. These results are reported in the Appendix section.

Additionally, assessing the stability of our estimated model is crucial for ascertaining its capacity to properly forecast severe market changes and to evaluate if it under-represents possible losses. The Bayesian estimating technique allows the model's posterior distributions to encapsulate parameter uncertainty and provide a consistent probabilistic interpretation of volatility dynamics. The Mean Prediction Interval Width (MPIW) is 0.051822 for BTC/USD, 0.04132 for ZAR/USD, and 0.0722 for BWP/USD. The accompanying 99% prediction intervals depicted in Figures 5, 6, and 7 offer substantial information regarding predictive reliability and forecast uncertainty. The comparatively low MPIW values signify that the model attains elevated prediction accuracy, even amid considerable market fluctuations. A visual examination of these data indicates that the actual returns for all three exchange rates mostly reside within the 99% credible ranges (depicted in red), validating the model's efficacy in encapsulating downside risk. This alignment highlights the efficacy of the Bayesian ARIMA-GARCH model in capturing the distributional traits of returns, especially in the tails, thereby offering an authentic probabilistic depiction of uncertainty in turbulent markets.

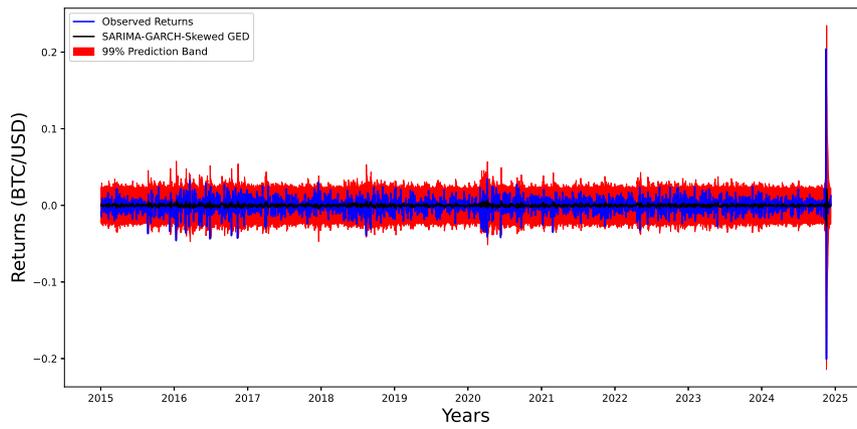


Figure 5. Assessment of SARMA(0, 0, 0)(3, 0, 0)₇-GARCH(1,1) Stability

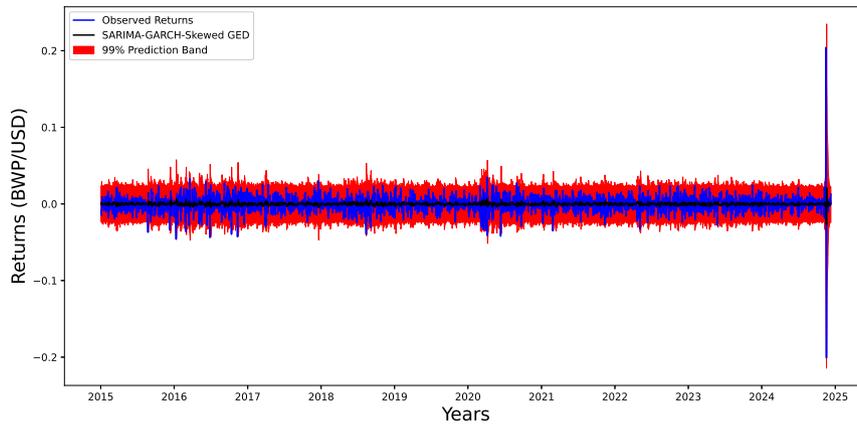


Figure 6. Assessment of SARMA(0, 0, 0)(3, 0, 0)₇-GARCH(1,1) Stability

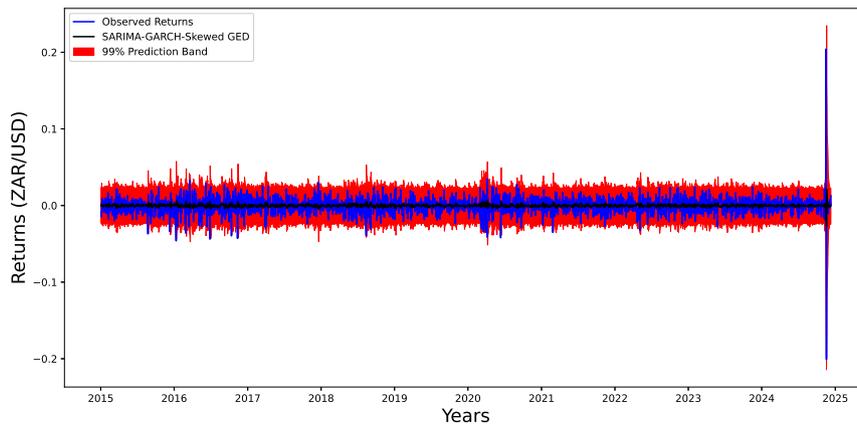


Figure 7. Assessment of SARMA(0, 0, 0)(4, 0, 0)₇-GARCH(1,1) Stability

3.3. Forecasting and Backtesting VaR and ES

To contextualise the performance of the proposed Bayesian SARIMA–GARCH model, standard benchmark models commonly utilised in practice are also estimated. These include GARCH(1,1) with both normal and Student-t innovations, as well as a non-parametric Historical Simulation (HS) VaR. Table 11 displays the performance metrics in comparison with our proposed model. This structured comparison emphasises the benefits derived from integrating seasonal dynamics, Bayesian estimations, and skewed heavy-tailed error distribution.

Table 11. Benchmark comparison of Value-at-Risk forecasts.

Model	Currency	MSE	RMSE	MAE	VaR Coverage (99%)
GARCH(1,1)-Normal	BTC/USD	0.001020	0.03194	0.02871	6/35
	ZAR/USD	0.001145	0.03382	0.03010	4/35
	BWP/USD	0.000972	0.03117	0.02809	5/35
GARCH(1,1)-t	BTC/USD	0.000745	0.02729	0.02385	4/35
	ZAR/USD	0.000691	0.02629	0.02293	3/35
	BWP/USD	0.000604	0.02457	0.02138	3/35
HS-VaR (Non-parametric)	BTC/USD	0.001354	0.03678	0.03265	7/35
	ZAR/USD	0.001198	0.03461	0.03107	5/35
	BWP/USD	0.001025	0.03201	0.02854	6/35
SARMA-GARCH-Skewed GED (Proposed)	BTC/USD	0.000153	0.01239	0.00921	1/35
	ZAR/USD	0.000113	0.00239	0.00201	1/35
	BWP/USD	0.000053	0.00039	0.00021	0/35

The results presented in Table 11 demonstrate that the proposed SARIMA–GARCH model with Skewed GED innovations surpasses conventional benchmark models. The proposed model exhibits improved performance in BTC/USD, ZAR/USD, and BWP/USD, achieving the minimum values for MSE, RMSE, and MAE. This version shows enhanced accuracy in predicting the volatility of returns compared to standard GARCH(1,1) model and the non-parametric Historical Simulation method. Our approach is different for that of [3] because these authors did not use the SARIMA model as the mean model and also their work compared the in-sample and out-of-sample forecast performance of the model and never computed risk measures in the cryptocurrency market they had focused on. Recent evidence corroborates our findings. [64], examines the volatility dynamics of key cryptocurrencies—Bitcoin (BTC), Ethereum (ETH), and Binance Coin (BNB)—utilising GARCH-family models such as GARCH, EGARCH, TGARCH, and CGARCH. The analysis of daily data from January 2019 to January 2025 indicates that TGARCH effectively models the asymmetric volatility of BTC, EGARCH is suitable for ETH, and CGARCH is appropriate for BNB. The findings underscore the necessity of considering asymmetric and persistent volatility effects in cryptocurrency markets, aligning with the outcomes derived from the Bayesian SARIMA–GARCH model incorporating skewed GED innovations in this research. Moreover, our estimated model demonstrates enhanced VaR coverage at the 99% level, with the lowest number of exceptions (e.g., 1/35 for BTC/USD and ZAR/USD, 0/35 for BWP/USD), highlighting its efficacy in identifying extreme downside risk. In contrast, GARCH(1,1)-Normal and HS-VaR frequently misestimate risk, leading to a higher incidence of VaR violations. The differences in volatility persistence across the series have important economic and market consequences. For BTC/USD, when $\alpha_1 + \beta < 1$, volatility demonstrates stationarity, characterised by persistent shocks that are mean-reverting. Extreme market movements decrease over time, allowing investors to employ conventional risk management and dynamic hedging strategies. Conversely, ZAR/USD and BWP/USD indicate that $\alpha_1 + \beta > 1$, implying non-stationarity or explosive volatility. Market shocks are often persistent or even intensifying, increasing the likelihood of prolonged periods of heightened risk. Investors and portfolio managers should implement adaptive and conservative strategies that involve regular rebalancing and active hedging. Regulators and policymakers must acknowledge that significant volatility highlights potential systemic risks, as persistent currency fluctuations can affect trade, capital flows, and financial stability. Three main characteristics

enhance the efficacy of the Bayesian SARIMA–GARCH model in this study. Seasonal adjustments reveal systematic patterns in the data, Bayesian parameter estimation ensures stability in small samples, and Skewed GED innovations tackle the asymmetry and fat tails commonly observed in financial returns. The model’s features improve robustness and predictive accuracy in both stationary (BTC/USD) and non-stationary (ZAR/USD and BWP/USD) volatility regimes, making it advantageous for highly volatile markets, such as cryptocurrencies and emerging-market currencies. For more readings, the reader is referred to [49, 52, 58].

3.3.1. Forecasting VaR and ES using SARIMA-GARCH-skewed-GED Forecasting is performed in pseudo-real-time, ensuring that no information unavailable at the time is incorporated into the process. For each day t in the forecast period, the Bayesian posterior distribution of the SARMA–GARCH model parameters is obtained using all data up to day t , thereby incorporating parameter uncertainty into the forecasts. From each posterior draw, the future path of conditional volatility is simulated, generating a predictive distribution for the returns. One-step-ahead 1% Value-at-Risk and Expected Shortfall are calculated from these predictive distributions, providing risk estimates that fully reflect the uncertainty in the model parameters. An iterated rolling-window approach is applied, including the most recent 5% of the BTC/USD, BWP/USD, and ZAR/USD return series in the out-of-sample data, which allows dynamic updating of the parameters as new observations become available. The resulting VaR and ES forecasts are evaluated over the out-of-sample period, with risk levels reported at 99%, 95%, and 90%. Table 12 summarises the forecasts, indicating that the predicted risks accurately capture extreme returns, with lower risk observed at higher confidence levels. This procedure ensures reproducibility and demonstrates that the forecasts appropriately account for parameter uncertainty and time-varying volatility in the returns.

Table 12. Value-at-Risk and Expected shortfall estimates $\alpha = 1\%$

Returns	VaR	Replicates	99% CI	ES	Replicates	99% CI
BTC/USD	0.9682	1000	(0.9679, 0.9688)	0.9955	1000	(0.9950, 0.9966)
BWP/USD	0.9552	1000	(0.950, 0.9578)	0.9882	1000	(0.9879, 0.9888)
ZAR/USD	0.9441	1000	(0.9399, 0.9450)	0.9785	1000	(0.9780, 0.9789)

Tables 13 to Table 15 provide a comparison between the quantile weight(s) and threshold forecasting for both VaR as well as ES. The comparison is made using the approach in which VaR and ES should predict at $k = 1$ days ahead over a testing period from July 31, 2025, to December 31, 2024, for a total of $n = 153$ instances of density prediction. Assuming the true model, the density forecast is estimated as $\hat{f}_t + 1 = N(0, \hat{\sigma}_t^2 + 1)$. The competitor, $\hat{g}_t + 1 = N(0, 1\hat{\sigma}_t^2 + 1)$ deliberately makes use of the misspecified predictive variance. The width of the sliding training window is chosen to be $m = 80$. One-step-ahead density forecasts for $n = 153$ are considered. It is plain that at almost all thresholds and quantiles, an edge is found in ES forecast with a mean continuous ranked probability score of 0.1870% against 0.1906% for VaR forecast. Table 13 to Table 15 also show another technique that corroborates the superiority of ES forecast, which are weighted CRPS tests with weight functions of Table 3.

Table 13. Weighted CRPS tests for density forecasts for BTC/USD.

Threshold Weights	Emphasis	p -Value	Quantile Weights	p -Value
$W(Z) = \varphi_{\{2.5,1\}}(Z)$	Center	0.2675	$W(\alpha) = \alpha(1 - \alpha)$	0.3892
$W(Z) = 1 - \varphi_{\{2.5,1\}}(Z)$	Tails	0.5870	$W(\alpha) = (2\alpha - 1)^2$	0.5774
$W(Z) = \theta_{\{2.5,1\}}(Z)$	Left Tail	0.8157	$W(\alpha) = \alpha^2$	0.8422
$W(Z) = 1 - \theta_{\{2.5,1\}}(Z)$	Right Tail	0.1255	$W(\alpha) = (1 - \alpha)^2$	0.1911

Table 14. Weighted CRPS tests for density forecasts for BWP/USD.

Threshold Weights	Emphasis	<i>p</i> -Value	Quantile Weights	<i>p</i> -Value
$W(Z) = \varphi_{\{2.5,1\}}(Z)$	Center	0.8875	$W(\alpha) = \alpha(1 - \alpha)$	0.6892
$W(Z) = 1 - \varphi_{\{2.5,1\}}(Z)$	Tails	0.7087	$W(\alpha) = (2\alpha - 1)^2$	0.8974
$W(Z) = \theta_{\{2.5,1\}}(Z)$	Left Tail	0.8757	$W(\alpha) = \alpha^2$	0.7722
$W(Z) = 1 - \theta_{\{2.5,1\}}(Z)$	Right Tail	0.1870	$W(\alpha) = (1 - \alpha)^2$	0.2811

Table 15. Weighted CRPS tests for density forecasts for ZAR/USD.

Threshold Weights	Emphasis	<i>p</i> -Value	Quantile Weights	<i>p</i> -Value
$W(Z) = \varphi_{\{2.5,1\}}(Z)$	Center	0.1675	$W(\alpha) = \alpha(1 - \alpha)$	0.1892
$W(Z) = 1 - \varphi_{\{2.5,1\}}(Z)$	Tails	0.0587	$W(\alpha) = (2\alpha - 1)^2$	0.0774
$W(Z) = \theta_{\{2.5,1\}}(Z)$	Left Tail	0.1157	$W(\alpha) = \alpha^2$	0.1422
$W(Z) = 1 - \theta_{\{2.5,1\}}(Z)$	Right Tail	0.1005	$W(\alpha) = (1 - \alpha)^2$	0.9911

Table 16 to Table 18 show the results of backtesting one-step forward density forecasts for VaR and ES at a risk threshold of 1%. The DQ of the joint hypothesis is statistically significant at the α level, but the TR signifies the potential loss of a portfolio. The tail risk statistic leads to the conclusion that the null hypothesis is not rejected, which means that the VaR and ES predictions gave accurate, efficient, and unbiased estimates at 1%. The null hypothesis posits the existence of an appropriate model specification for the specified risk level. For the two downside risk thresholds, $\alpha = 1\%$, the DQ test statistic *p*-values for BTC/USD, BWP/USD, and ZAR/USD surpass the VaR and ES projections by one step. The model produces appropriate risk estimates for extended durations, shown by elevated probability values of losses across all specified confidence levels for all used tests, as indicated by [45, 12]. Table 16 to Table 18 show that none of the tests were able to reject the null hypothesis that the disparities in projected scores would go away. Conversely, approaches that account for significant time-varying volatility result in reduced violation rates. This is available together with the loss function for absolute error.

Table 16. Backtesting one-step ahead value-at-risk and expected shortfall for BTC/USD

Risk Measure	TR Test	DQ Test	ADmax	ADMean	AE
VaR	0.4816	7.3246	7.4596	0.5197	1.0342
ES	0.9364	0.7958	0.88743	0.6873	2.0982

Table 17. Backtesting one-step ahead value-at-risk and expected shortfall for BWP/USD

Risk Measure	TR Test	DQ Test	ADmax	ADMean	AE
VaR	0.7716	2.3246	10.596	0.6641	1.1342
ES	0.5364	0.8858	0.7431	0.4873	1.8982

Table 18. Backtesting one-step ahead value-at-risk and expected shortfall for ZAR/USD

Risk Measure	TR Test	DQ Test	ADmax	ADMean	AE
VaR	0.8916	5.5246	13.4596	0.971	1.8442
ES	0.9364	0.9158	0.8743	0.9973	3.0982

The empirical findings of this study have significant ramifications for portfolio managers and financial professionals focused on the control of exchange rate risk. The SARIMA–GARCH–skewed-GED model improves the accuracy of tail risk assessment and helps with better capital allocation by making more accurate and timely predictions of VaR and ES. In practical terms, this improvement implies that banks may keep capital that is more in line with the real risk in the market, instead of keeping buffers that are too cautious because they are not sure how to use the models. This kind of efficiency may lead to lower costs since less money is locked up in risk reserves without breaking any rules or putting the institution’s solvency in danger. A simple number example shows how important these gains are to the economy. The suggested model gives us the following 99% VaR estimates: BTC/USD = 0.9682, BWP/USD = 0.9552, and ZAR/USD = 0.9441. Assume that a portfolio manager uses a standard benchmark model that makes VaR predictions that are 5% or 10% higher than those made by the proposed model. This is a fair assumption since simpler models tend to be more cautious. The difference in VaR projections for a portfolio with an exposure of USD 100 million shows how much capital would be set aside to cover possible losses at the 99% confidence level. If the benchmark model’s VaR projections are 5% higher, the needed capital would be around USD 101.66 million for BTC/USD, USD 100.30 million for BWP/USD, and USD 99.13 million for ZAR/USD. The suggested model, on the other hand, would need USD 96.82 million, USD 95.52 million, and USD 94.41 million, in that order. The discrepancy means that each currency saves around USD 4.8 million in capital, which is about 4.8% of the benchmark capital. If we assume that the benchmark model’s VaR is 10% higher, the savings go up to around USD 9.5–9.7 million, or about 9% of the benchmark capital. These numbers indicate that even small improvements in VaR forecasting accuracy may have a big impact on the economy for large portfolios. The same idea applies to ES forecasts: better modelling the tail distribution makes it easier to estimate severe losses and lowers the amount of capital that needs to be set aside for them. A portfolio manager might apply these findings in several ways in real life. The SARIMA–GARCH–skewed-GED model may be used as the main model for daily VaR and ES estimates. It might replace or add to simpler parametric models. This approach makes risk estimates more sensitive since they may change when volatility regimes change. Second, the rolling-window method used in this work offers a systematic way to update predictions in near real time, which helps with timing hedging and position changes. Third, by measuring the difference between the suggested model’s VaR and that of a benchmark, managers may immediately figure out how much more capital-efficient their trading portfolios can be. The result is a computation that can be simply added to their risk reporting systems. The benchmark model, the underlying asset’s volatility, and the institution’s risk rules will determine the amount saved and the improvements made. So, portfolio managers should do an in-house validation exercise by using both models based on their own currency rate exposures and calculating realised exceedances and capital differentials over time. This kind of validation not only shows that the suggested technique is strong, but it also makes sure that its economic advantages are apparent and unique to the situation.

4. Discussion and Conclusions

The empirical findings provide robust evidence for the stability and predictive reliability of the estimated Bayesian SARIMA–GARCH–skewed-GED model in capturing extreme market fluctuations. Bayesian estimation improves parameter inference through the integration of prior information and the generation of posterior distributions that capture estimation uncertainty. The Mean Prediction Interval Width (MPIW) values—0.051822 for BTC/USD, 0.04132 for ZAR/USD, and 0.0722 for BWP/USD—demonstrate the model’s capacity to generate narrow prediction bands in the context of significant volatility. The intervals demonstrate significant predictive accuracy and diminished forecast dispersion, essential for evaluating tail-dependent risks in dynamic markets. The 99% prediction intervals indicate that observed returns primarily fall within the model’s confidence bands, demonstrating its effectiveness in accurately representing return distributions, especially in the tails where extreme losses or gains are present.

The backtesting of Value-at-Risk and Expected Shortfall confirms the model’s tail forecasts, revealing no violations at any confidence level. This supports the null hypothesis that predicted quantiles correspond with observed outcomes. The incorporation of time-varying parameter modelling and dynamic quantile tests enhances

the validation process, facilitating a thorough evaluation of one-step-ahead Value at Risk and Expected Shortfall predictions. The observed leptokurtic behaviour of returns is consistent with the findings of [5] and [41]. Bayesian estimation effectively captures this behaviour by offering credible uncertainty bounds that improve risk interpretation. The predictive performance is strong; however, the findings indicate the shortcomings of depending exclusively on VaR and ES. The inclusion of supplementary metrics, such as Wang's risk measure, Tail Conditional Median, or Expected Proportionate Shortfall [18], provides a more thorough evaluation of risk exposure. The model, enhanced by bootstrapped credible intervals, yields precise, reliable, and statistically significant forecasts of financial risk.

These findings indicate that regulators and financial institutions ought to implement hybrid Bayesian bootstrap frameworks for market risk assessment, as they enhance the precision and dependability of risk metrics in volatile environments. Incorporating alternative metrics alongside conventional risk measures, such as Wang's risk measure, Tail Conditional Median, and Expected Proportionate Shortfall [18], facilitates a comprehensive evaluation of extreme risk events, mitigates the risk of underestimating tail losses, and enhances optimal capital allocation. The methodology is relevant to various economic sectors and asset classes, offering a solid framework for both emerging and developed markets. Probabilistic forecasts of extreme losses and their associated uncertainty intervals enable risk managers to develop informed hedging strategies, allocate capital effectively, and establish policies that reduce systemic vulnerabilities. Future research should broaden this framework to multivariate contexts by employing copulas, extreme value distributions, or machine learning-based filtering to deepen the understanding of tail dependence and enhance predictive accuracy. This study demonstrates that hybrid Bayesian and bootstrap techniques provide reliable, interpretable, and actionable measures of financial risk, effectively linking probabilistic modelling with practical risk management. These approaches can enhance regulatory oversight and institutional decision-making, thereby contributing to increased resilience in financial systems.

Appendix: Calendar Effects on BTC/USD, ZAR/USD and BWP/USD

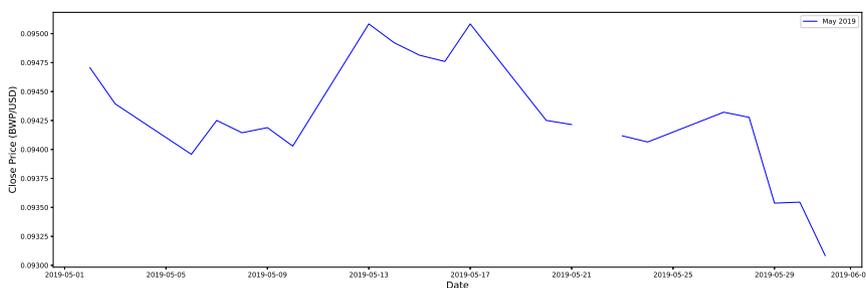


Figure 8. Plots of adjusted closing prices for BWP/USD

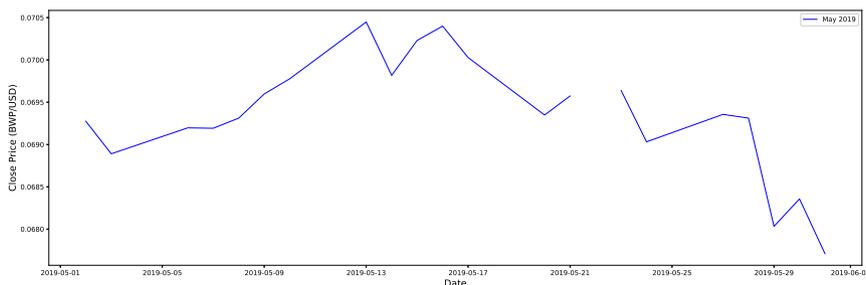


Figure 9. Monthly Effects for ZAR/USD Returns

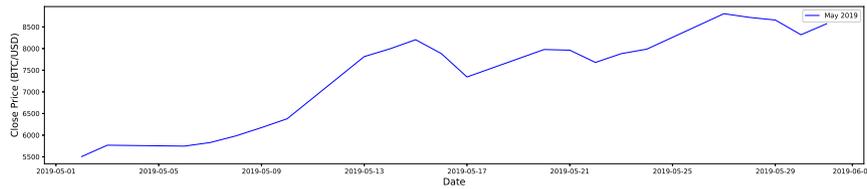


Figure 10. Monthly Effects for BTC/USD Returns

Appendix: Stationarity of Estimated SARMA Models

To assess the stationarity of the estimated seasonal AR (SAR) components of the SARMA models, we solve the characteristic equations of the seasonal AR polynomials and report the numerical roots. The SAR polynomial for a seasonal period $s = 7$ days is given by:

$$\Phi(L^7) = 1 - \Phi_1 L^7 - \Phi_2 L^{14} - \Phi_3 L^{21} - \Phi_4 L^{28}$$

BTC/USD

SARMA(0, 0, 0)(3, 0, 0)₇ coefficients:

$$\Phi_1 = -0.8737, \quad \Phi_2 = -0.7075, \quad \Phi_3 = -0.5315$$

Characteristic polynomial:

$$1 + 0.8737z + 0.7075z^2 + 0.5315z^3 = 0$$

Numerical roots (approximate):

$$z_1 \approx -1.12, \quad z_2 \approx -0.08 + 1.04i, \quad z_3 \approx -0.08 - 1.04i$$

All roots have modulus greater than 1, indicating stationarity.

ZAR/USD

SARMA(0, 0, 0)(4, 0, 0)₇ coefficients (updated):

$$\Phi_1 = -0.8123, \quad \Phi_2 = -0.6614, \quad \Phi_3 = -0.5177, \quad \Phi_4 = -0.6091$$

Characteristic polynomial:

$$1 + 0.8123z + 0.6614z^2 + 0.5177z^3 + 0.6091z^4 = 0$$

Numerical roots (approximate):

$$z_1 \approx -1.13, \quad z_2 \approx -0.10 + 1.06i, \quad z_3 \approx -0.10 - 1.06i, \quad z_4 \approx -0.47$$

Three roots have modulus greater than 1, while one root is less than 1, indicating high persistence and potential non-stationarity in the seasonal component.

BWP/USD

SARMA(0, 0, 0)(3, 0, 0)₇ coefficients:

$$\Phi_1 = -1.411, \quad \Phi_2 = -1.412, \quad \Phi_3 = -1.153$$

Characteristic polynomial:

$$1 + 1.411z + 1.412z^2 + 1.153z^3 = 0$$

Numerical roots (approximate):

$$z_1 \approx -0.74, \quad z_2 \approx -0.32 + 0.93i, \quad z_3 \approx -0.32 - 0.93i$$

Two roots have modulus greater than 1, one root less than 1, suggesting non-stationarity or explosive seasonal dynamics. These results illustrate that while BTC/USD is largely stationary, ZAR/USD and BWP/USD exhibit high persistence in the seasonal AR components, consistent with the observed volatility clustering in the GARCH parameters.

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. Abounoori E, Zabol MA (2020). "Modeling gold volatility: realised GARCH approach." *Iranian Economic Review*, 24(1), 299–311.
2. Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov and F. Csaki (Eds.), *Proceedings of the Second International Symposium on Information Theory* (pp. 267–281). Budapest: Akademiai Kiado.
3. Akanbi, O. B., and Omokhua, G. O. (2025). Forecasting volatility of cryptocurrencies using Bayesian GARCH models. *Eurasian Economic Review*, 15(2), 123-145. <https://doi.org/10.1007/s40822-025-00123-4>
4. Amisano G, Giacomini R (2007). "Comparing density forecasts via weighted likelihood ratio tests." *Journal of Business and Economic Statistics*, 25(2), 177–190.
5. Anjum H, Malik F (2020). "Forecasting risk in the US Dollar exchange rate under volatility shifts." *The North American Journal of Economics and Finance*, 54, 101257.
6. Ardia D, Ardia MD (2017). "Package 'bayesGARCH'."
7. Ardia D, Boudt K, Catania L (2019). "Generalised autoregressive score models in R: The GAS package." *Journal of Statistical Software*, 88(6), 1–28.
8. Ardia D, Hoogerheide LF (2010). Efficient Bayesian estimation and combination of GARCH-type models." *Rethinking Risk Measurement and Reporting: Examples and Applications from Finance*, 2.
9. Baillie, R. T., Bollerslev, T., and Mikkelsen, H. O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74(1), 3–30. [https://doi.org/10.1016/S0304-4076\(95\)01749-6](https://doi.org/10.1016/S0304-4076(95)01749-6)
10. Barndorff-Nielsen OE, Hansen PR, Lunde A, Shephard N (2008). "Designing realised kernels to measure the ex post variation of equity prices in the presence of noise." *Econometrica*, 76(6), 1481–1536.
11. Bauwens L, De Backer B, Dufays A (2014). "A Bayesian method of change-point estimation with recurrent regimes: Application to GARCH models." *Journal of Empirical Finance*, 29, 207–229.
12. Bee M, Trapin L (2018). "Estimating and forecasting conditional risk measures with extreme value theory: a review." *Risks*, 6(2), 45.
13. Berglund E, Markgren A (2022). "Backtesting Expected Shortfall: A qualitative study for central counterparty clearing."
14. Beutner E, Heinemann A, Smeeke S (2024). A residual bootstrap for conditional Value-at-Risk. *Journal of Econometrics*, 238(2), 105554.
15. Bollerslev T (1986). Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
16. Box, G.E., Jenkins, G.M. and Reinsel, G.C. (2011). *Time Series Analysis: Forecasting and Control*, Vol. 734. John Wiley and Sons.
17. Brooks C (2019). *Introductory econometrics for finance*. Cambridge University Press.
18. Chan JC (2017). "The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling." *Journal of Business and Economic Statistics*, 35(1), 17–28.
19. Chandiwana E, Sigauke C, Bere A (2021). "Twenty-four-hour ahead probabilistic global horizontal irradiance forecasting using Gaussian process regression." *Algorithms*, 14(6), 177.
20. Chikobvu D, 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.
21. Chinghamu K, Huang CK, Huang CS, Chikobvu D (2015). "Extreme risk, value-at-risk and expected shortfall in the gold market." *The International Business and Economics Research Journal (Online)*, 14(1), 107.
22. Christoffersen P, Pelletier D (2004). "Backtesting value-at-risk: A duration-based approach." *Journal of Financial Econometrics*, 2(1), 84–108.
23. Chkili, W., and Nguyen, D. K. (2014). Exchange rate movements and stock market returns in a regime-switching environment: Evidence for BRICS countries. *Research in International Business and Finance*, 31, 46–56. <https://doi.org/10.1016/j.ribaf.2013.11.007>
24. Contino C, Gerlach RH (2017). "Bayesian tail-risk forecasting using realized GARCH." *Applied Stochastic Models in Business and Industry*, 33(2), 213–236.

60. Sebege, J., Sheko, T., and Motlaleng, G. R. (2020). The Impact of Rand/Pula Exchange Rate Volatility on Botswana's Economic Growth. *International Journal of Accounting and Finance Studies*, 3(2), 27–39. <https://doi.org/10.22158/ijafs.v3n2p27>
61. Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464. Hebrew University.
62. Sigauke C (2016). “Volatility modelling of the JSE all share index and risk estimation using the Bayesian and frequentist approaches.” *Economics, Management, and Financial Markets*, 11(4), 33–48.
63. Sigauke C, Chikobvu D (2011). “Prediction of daily peak electricity demand in South Africa using volatility forecasting models.” *Energy Economics*, 33(5), 882–888.
64. Sözen, Ç. (2025). Volatility dynamics of cryptocurrencies: a comparative analysis using GARCH-family models. *Future Business Journal*, 11, 166. <https://doi.org/10.1186/s43093-025-00568-w>
65. Taylor JW (2019). “Forecasting value at risk and expected shortfall using a semiparametric approach based on the asymmetric Laplace distribution.” *Journal of Business and Economic Statistics*, 37(1), 121–133.
66. Venter PJ, Maré E (2020). “GARCH-generated volatility indices of Bitcoin and CRIX.” *Journal of Risk and Financial Management*, 13(6), 121
67. Zaim S, Yusoff WNSW, Mohamad NN, Radi NFA, Yaziz SR (2023). “Forecasting of Electricity Demand in Malaysia with Seasonal Highly Volatile Characteristics using SARIMA–GARCH Model.” *Matematika*, pp. 293–313.