Journal of Statistical Models and Methods (2025)
1(1), 1–26, https://doi.org/10.65905/wvgkn788

Article Type: Research article

Forecasting Major GBP Currency Pairs Using Bayesian Methods

Talal Kurdi[1]

[1]Department of Mathematics, University of Manchester, Manchester, M13 9PL, UK

Talal Kurdi
talal.kurdi-2@postgrad.manchester.ac.uk

Article History

Received: 11 November 2025
Revised: 29 November 2025
Accepted: 06 December 2025

Keywords
Bayesian inference;
Skewed student-t distribution;
Exchange rate forecasting;
MCMC methods;
Currency pairs;
Tail risk

Abstract

This study develops a comprehensive Bayesian framework for modeling and forecasting daily log-returns of three major British Pound currency pairs: GBP/USD, GBP/EUR, and GBP/JPY. Using approximately 23 years of daily data, employing model selection criteria to identify the skewed Student-t distribution as the optimal specification among seven candidate distributions. Parameter estimation is conducted via an adaptive Metropolis-Hastings algorithm. The empirical results reveal several important characteristics of GBP currency returns. All three pairs exhibit substantially heavy tails, with posterior mean degrees of freedom ranging from 3.40 (GBP/JPY) to 4.84 (GBP/USD), indicating extreme movements occur far more frequently than predicted by normal distributions. Negative skewness parameters confirm asymmetric left-tail risk across all pairs. Daily volatility estimates range from 0.36% (GBP/EUR) to 0.51% (GBP/JPY). Comprehensive MCMC diagnostics confirm excellent convergence, with effective sample sizes exceeding 300 and acceptance rates in the optimal 22-24% range. The Bayesian framework provides honest uncertainty quantification, essential for risk management applications such as Value-at-Risk computation, while acknowledging the fundamental unpredictability of exchange rate movements.

1 Introduction

The foreign exchange market is the world's largest and most liquid financial market, with currency exchange rates playing a critical role in international trade, investment decisions, and economic policy, making accurate modeling and forecasting of these rates essential for financial institutions, corporations, and policymakers. This study aims to estimate the unknown parameters of the log-return model for major GBP currency pairs using Bayesian techniques. We aim to forecast British Pound vs. US Dollar (GBP/USD), British Pound vs. Euro (GBP/EUR), and British Pound vs. Japanese Yen (GBP/JPY). At first, suitable models are selected and tested to determine the best model based on known criteria. Then, the parameters of the best model are estimated using the Bayesian method via the Metropolis-Hastings algorithm. Finally, the posterior distribution is used to generate price forecasts.

1.1 Background and Motivation

The analysis of historical and future asset prices has captured the interest of many generations of researchers. Instead of studying the price P_t at intervals $t = 1, \ldots, T$ directly, the primary focus is on the log-returns defined by $R_t = \log(P_t) - \log(P_{t-1})$ for $t = 2, \ldots, T$ for several reasons (see Fergusson

Copyright © 2025 by the author. Published under Creative Commons Attribution (CC BY) license.

and Platen[1] and Platen and Rendek [2]). Therefore, many studies have been conducted to estimate the distributions of log-returns for financial assets, see, for example, Markowitz and Usmen [3] and Markowitz and Usmen [4]. Numerous studies have examined methods for estimating and forecasting currency exchange rates. Hsieh [5] studied the statistical properties of five daily foreign exchange rates in terms of the USD from 1974 to 1983. Nag and Mitra [6] used the neural network to forecast daily exchange rates of the German Deutsche Mark/USD, USD/GBP and JPY/USD. Moreover, Darvas and Schepp [7] forecasted the GBP/USD exchange rate. While numerous studies have examined currency exchange rate forecasting, the application of fully Bayesian methods with MCMC techniques remains relatively underutilized for GBP currency pairs. This study fills this gap by developing a comprehensive Bayesian framework for modeling and forecasting GBP/USD, GBP/EUR, and GBP/JPY exchange rates. The Bayesian paradigm naturally accommodates parameter uncertainty and provides probabilistic forecasts rather than point predictions, making it particularly suitable for financial decision-making under uncertainty.

1.2 Literature on Bayesian Modeling of Heavy-Tailed Returns

The application of Bayesian methods to heavy-tailed financial return distributions has received growing attention, driven by recognition that Gaussian assumptions systematically underestimate tail risk. The choice of prior distributions for scale and tail parameters poses unique methodological challenges, particularly in heavy-tailed models, where standard noninformative priors may yield improper posteriors. Ramos et al. [8] demonstrate that objective priors for a general family of distributions exhibit power-law behavior, with exponents determining whether the resulting posterior distributions are proper or improper. Their framework provides principled guidance for prior construction in models with heavy tails, showing that understanding the asymptotic behavior of priors is essential for ensuring valid Bayesian inference. This insight directly informs our specification of weakly informative priors: we adopt proper priors for all parameters (including the degrees of freedom ν and scale σ) that impose minimal constraints while avoiding the propriety issues that can arise with purely noninformative specifications. The use of Student-t distributions for modeling financial returns has a well-established history in econometrics, with extensions to skewed variants capturing the asymmetric tail behavior commonly observed in asset returns. Our contribution extends this literature by providing comprehensive empirical evidence for three major GBP currency pairs and, through rigorous model comparison (AIC/BIC across seven candidate distributions), demonstrating that the skewed Student-t specification dominates alternatives. Moreover, we explicitly incorporate parameter uncertainty into forecast distributions through full Bayesian inference, showing that credible intervals achieve proper empirical coverage. At the same time, approaches conditioning on point estimates would systematically understate forecast uncertainty.

The remainder of this study is organized as follows: Section 2 discusses the data selection procedure and shows the descriptive statistics of the collected data. Section 3 demonstrates how the model is selected. The theoretical illustration of Bayesian analysis is shown in Section 4. The empirical part of the Bayesian method is presented in Section 5. Moreover, Bayesian price forecasting is shown in Section 6. Finally, Section 7 presents the study's conclusion.

2 Data Collection and Descriptive Statistics

Data was collected for daily trading of three pairs GBP/USD, GBP/EUR and GBP/JPY from 01/12/2003 to 10/10/2025. Table 1 presents the descriptive statistics for the three GBP currency pairs over the study period. All three series exhibit positive skewness, with GBP/EUR showing the highest value at 0.9314, indicating right-tailed distributions characteristic of price levels with occasional large upward movements. However, this positive skewness in price levels does not directly translate into the distributional properties of log-returns, which are our focus for modeling. The reported kurtosis values (2.17 to 2.38) represent excess kurtosis (kurtosis minus 3), indicating that the actual kurtosis values range from 5.17 to 5.38, which exceed the normal distribution benchmark of 3. These elevated kurtosis values confirm the presence of heavier tails and more extreme values than would be expected under normality—a well-documented stylized fact in exchange rate data. Notably, GBP/JPY exhibits substantially higher volatility (SD =

30.96) than GBP/USD (SD = 0.24) and GBP/EUR (SD = 0.13), reflecting the Japanese Yen's particular sensitivity to global risk sentiment and carry trade dynamics. Table 2 shows the descriptive statistics for the log-returns of GBP currency pairs for 23 years. All three series exhibit negative skewness, indicating asymmetric distributions with heavier left tails. While kurtosis is not reported in Table 2, the substantially elevated excess kurtosis values computed from the data (GBP/USD: 11.82, GBP/EUR: 8.90, GBP/JPY: 16.89) correspond to actual kurtosis values of 14.82, 11.90, and 19.89 respectively, dramatically exceeding the normal distribution benchmark of 3. These values reflect the presence of extreme events and heavy-tailed behavior characteristic of foreign exchange returns. The combination of negative skewness and excess kurtosis motivates the use of the skewed Student-t distribution, which accommodates both asymmetry and heavy tails.

Figures 1-3 illustrate the temporal evolution of the three currency pairs from 2003 to 2025, with vertical dashed lines marking three major economic events: the 2008 financial crisis, the COVID-19 pandemic, and Brexit. Several important patterns emerge from these visualizations. First, all three pairs exhibit pronounced responses to the marked crisis periods, with the 2008 financial crisis triggering the most severe disruptions. GBP/USD declined sharply from above 2.0 to approximately 1.4, while GBP/JPY collapsed from its peak near 250 to below 120. Second, a clear structural break is evident following the Brexit referendum in 2016, particularly in GBP/EUR and GBP/USD, which settled into persistently lower ranges and never recovered to their pre-crisis levels. Third, all series display volatility clustering, with periods of high turbulence concentrated around major economic events followed by relatively calm intervals. This heteroskedastic behavior, combined with the observed skewness and excess kurtosis, provides strong empirical justification for employing flexible parametric distributions such as the skewed Student-t in our Bayesian modeling framework. Finally, GBP/JPY exhibits more cyclical behavior compared to the other pairs, with a notable recovery trend emerging since 2020, potentially reflecting changing global monetary policy dynamics and risk appetite.

Table 1: Descriptive statistics of currency exchange rate price levels, 2003-2025.

Statistic	GBP/USD	GBP/EUR	GBP/JPY
Mean	1.5199	1.2391	167.0510
Median	1.5152	1.1828	158.3675
SD	0.2389	0.1271	30.9564
Min	1.0728	1.0157	117.1800
Max	2.1082	1.5262	250.2870
Q1	1.3036	1.1480	141.3745
Q3	1.6599	1.3076	193.6970
Skewness	0.5005	0.9314	0.5034
Kurtosis	2.1685	2.3839	2.2516

Note: Kurtosis values reported are excess kurtosis (kurtosis - 3). Actual kurtosis values are 5.17, 5.38, and 5.25 for GBP/USD, GBP/EUR, and GBP/JPY, respectively.

Table 2: Descriptive statistics for daily log-returns of currency exchange rates, 2003-2025.

Series	Sample Size	Mean	SD	Skewness
GBP/USD	5,685	-4.45×10^{-5}	0.00589	−0.670
GBP/EUR	5,726	-3.77×10^{-5}	0.00493	−0.295
GBP/JPY	5,689	1.36×10^{-5}	0.00775	−0.709

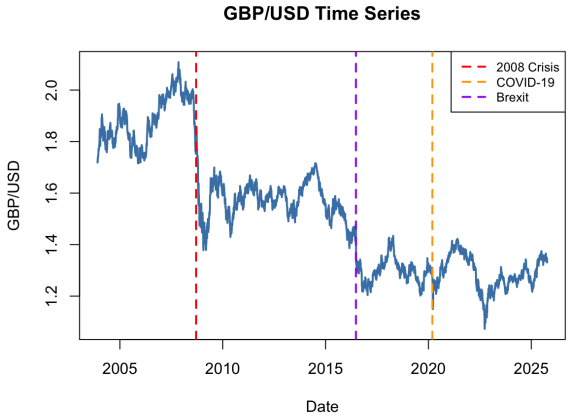

Figure 1: Time series plot for GBP/USD exchange rates, 2003-2025. Vertical dashed lines indicate the 2008 financial crisis, Brexit referendum (2016), and COVID-19 pandemic (2020).

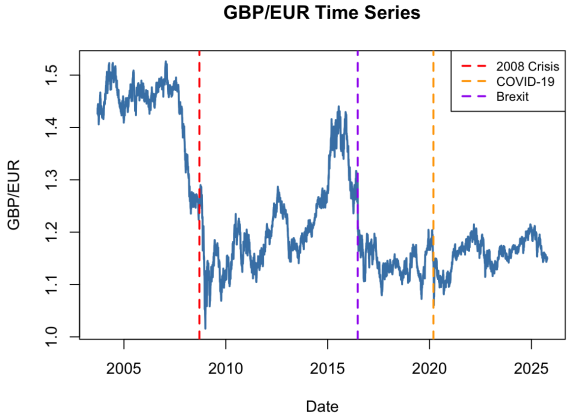

Figure 2: Time series plot for GBP/EUR exchange rates, 2003-2025. Vertical dashed lines indicate the 2008 financial crisis, Brexit referendum (2016), and COVID-19 pandemic (2020).

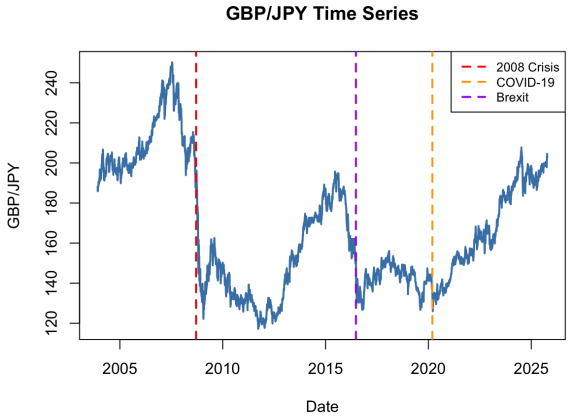

Figure 3: Time series plot for GBP/JPY exchange rates, 2003-2025. Vertical dashed lines indicate the 2008 financial crisis, Brexit referendum (2016), and COVID-19 pandemic (2020).

3 Model Selection

3.1 Probability Distributions

To model the log-returns of currency exchange rates, we consider several probability distributions commonly used in financial econometrics:

Normal Distribution: The standard Gaussian distribution with probability density function (pdf):

$$f(x \mid \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad x \in \mathbb{R},\ \mu \in \mathbb{R},\ \sigma > 0,$$

where μ is the mean and σ^2 is the variance. While widely used, it often underestimates the probability of extreme events in financial returns.

Student's t-Distribution: A heavy-tailed distribution with pdf:

$$f(x \mid \nu) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}, \quad x \in \mathbb{R},\ \nu > 0,$$

where ν is the degrees of freedom parameter. It better captures the excess kurtosis typically observed in financial returns, leading to more frequent extreme values than in the normal distribution.

Cauchy Distribution: An extremely heavy-tailed distribution with pdf:

$$f(x \mid x_0, \gamma) = \frac{1}{\pi\gamma\left[1 + \left(\frac{x-x_0}{\gamma}\right)^2\right]}, \quad x \in \mathbb{R},\ x_0 \in \mathbb{R},\ \gamma > 0,$$

where x_0 is the location parameter and γ is the scale parameter. It is a special case of Student's t with $\nu = 1$ and has no defined mean or variance, representing the most extreme tail behavior.

Logistic Distribution: Similar to the normal distribution but with slightly heavier tails, with pdf:

$$f(x \mid \mu, s) = \frac{e^{-\frac{x-\mu}{s}}}{s\left(1 + e^{-\frac{x-\mu}{s}}\right)^2}, \quad x \in \mathbb{R},\ \mu \in \mathbb{R},\ s > 0,$$

where μ is the location parameter and s is the scale parameter.

Laplace Distribution: Also known as the double exponential distribution, with pdf:

$$f(x \mid \mu, b) = \frac{1}{2b} \exp\left(-\frac{|x-\mu|}{b}\right), \quad x \in \mathbb{R},\ \mu \in \mathbb{R},\ b > 0,$$

where μ is the location parameter and b is the scale parameter. It has heavier tails than the normal distribution and is characterized by a sharp peak at the mode, making it suitable for modeling returns with a higher frequency of observations near zero.

Generalized Error Distribution (GED): A flexible distribution with pdf:

$$f(x \mid \mu, \alpha, \nu) = \frac{\nu \exp\left[-\left|\frac{x-\mu}{\alpha}\right|^\nu\right]}{2\alpha\Gamma(1/\nu)}, \quad x \in \mathbb{R},\ \mu \in \mathbb{R},\ \alpha > 0,\ \nu > 0,$$

where μ is the location, α is the scale, and ν is the shape parameter. It nests the normal distribution ($\nu = 2$) and Laplace distribution ($\nu = 1$) as special cases, allowing for various tail behaviors.

Skewed Normal Distribution: An extension of the normal distribution that incorporates a skewness parameter α, with pdf:

$$f(x \mid \mu, \sigma, \alpha) = \frac{2}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right)\Phi\left(\alpha\frac{x-\mu}{\sigma}\right), \quad x \in \mathbb{R},\ \mu \in \mathbb{R},\ \sigma > 0,\ \alpha \in \mathbb{R},$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal pdf and cdf, respectively. This allows for asymmetry in the distribution of returns.

Skewed Student's t-Distribution: Combines the heavy-tailed property of the Student's t-distribution with asymmetry. Its PDF is given by:

$$f(x \mid \mu, \sigma, \alpha, \nu) = \frac{2}{\sigma}t_\nu\left(\frac{x-\mu}{\sigma}\right)T_{\nu+1}\left(\alpha\frac{x-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+(x-\mu)^2/\sigma^2}}\right),$$

for $x \in \mathbb{R}$, $\mu \in \mathbb{R}$, $\sigma > 0$, $\alpha \in \mathbb{R}$, $\nu > 0$, where $t_\nu(\cdot)$ and $T_\nu(\cdot)$ are the Student's t PDF and CDF with ν degrees of freedom, μ is the location parameter, σ is the scale parameter, α is the shape parameter controlling skewness, and ν is the degrees of freedom parameter. This distribution flexibly captures both asymmetry and heavy tails characteristic of financial returns.

3.2 Model Selection Criteria

To compare the goodness-of-fit across different distributions, we employ two information criteria:

Akaike Information Criterion (AIC): Proposed by Akaike [9], the AIC balances model fit and complexity:

$$\text{AIC} = 2k - 2\ln(\hat{L})$$

where k is the number of parameters and $\hat{L}$ is the maximized likelihood function. Lower AIC values indicate better model performance.

Bayesian Information Criterion (BIC): Introduced by Schwarz [10], the BIC imposes a stronger penalty for model complexity:

$$\text{BIC} = k\ln(n) - 2\ln(\hat{L})$$

where n is the sample size. The BIC tends to favor more parsimonious models than the AIC, particularly in large samples. As with AIC, lower values indicate superior model fit.

Table 3 presents the model comparison results based on the AIC and BIC for seven candidate distributions fitted to the log-returns of the three GBP currency pairs. The skewed Student-t distribution consistently emerges as the best-fitting model across all three currency pairs, achieving the lowest AIC and BIC values in each case (GBP/USD: AIC = -42923.34, BIC = -42896.75; GBP/EUR: AIC = -45317.31, BIC = -45290.70; GBP/JPY: AIC = -40447.61, BIC = -40421.03). The standard Student-t distribution ranks second, with information criteria values only marginally higher than the skewed variant, suggesting that both heavy tails and asymmetry are important features of these currency return distributions. The normal and skewed normal distributions perform poorly across all pairs, with substantially higher AIC and BIC values, indicating their inadequacy in capturing the leptokurtic nature of exchange rate returns. Similarly, the Cauchy distribution, despite its heavy tails, yields the worst fit among all candidates, likely due to its overly extreme tail behavior and undefined moments. The superiority of the skewed Student-t distribution aligns with the positive skewness observed in Table 2 and confirms the presence of asymmetric risk in GBP currency returns. Figure 4 provides visual confirmation of the model selection results by displaying histograms of the log-returns overlaid with the fitted skewed Student-t density curves for each currency pair. The fitted distributions demonstrate excellent agreement with the empirical data across all three pairs. The smooth density curves closely track the histogram shapes, successfully capturing both the central concentration of returns and the tail behavior. Notably, the skewed Student-t

distribution accommodates the slight asymmetry visible in the empirical distributions, particularly evident in GBP/EUR where the right tail appears somewhat heavier than the left. The quality of fit is particularly impressive in the central region, where the majority of observations concentrate, while also providing adequate coverage of the tails where extreme movements occur. These visualizations corroborate the information criteria rankings in Table 3 and provide strong evidence that the skewed Student-t distribution is an appropriate choice for modeling GBP currency log-returns in the subsequent Bayesian estimation framework.

Table 3: Model comparison: AIC and BIC for log-returns distribution.

	GBP/USD		GBP/EUR		GBP/JPY	
Distribution	AIC	BIC	AIC	BIC	AIC	BIC
Skew Student T	**-42923.34**	**-42896.75**	**-45317.31**	**-45290.70**	**-40447.61**	**-40421.03**
Student T	-42861.90	-42841.97	-45288.61	-45268.65	-40432.62	-40412.68
Logistic	-42847.96	-42834.67	-45210.85	-45197.55	-40184.32	-40171.03
Laplace	-42797.28	-42783.99	-45259.11	-45245.80	-40346.80	-40333.50
Normal	-42249.81	-42236.52	-44579.49	-44566.19	-39143.96	-39130.66
Skew Normal	-42247.81	-42227.87	-44577.49	-44557.53	-39141.95	-39122.01
Cauchy	-41640.30	-41627.01	-44173.37	-44160.06	-39454.78	-39441.48

Note: Lower AIC and BIC values indicate better model fit. Bold values indicate the best-fitting model for each currency pair and criterion.

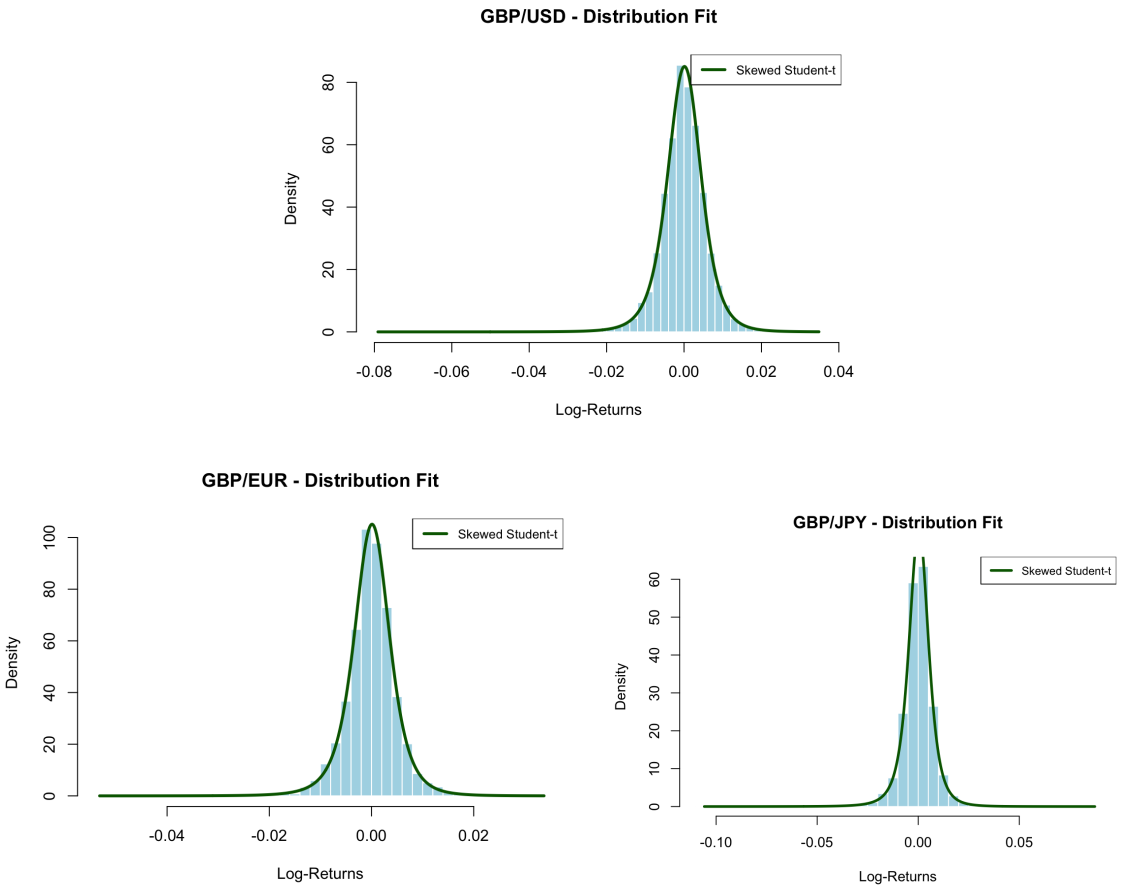

Figure 4: Fitted skewed Student-t distribution (curves) overlaid on histograms of daily log-returns for GBP/USD, GBP/EUR, and GBP/JPY, 2003-2025.

3.3 Serial Correlation Assessment

To validate the independence assumption underlying our model specification, we conducted Ljung-Box tests for serial correlation in the log-returns at multiple lag orders. Table 4 presents the test statistics and p-values for all three currency pairs at lags 1, 5, and 10. The null hypothesis is that there is no autocorrelation up to the specified lag. The results provide strong support for the independence assumption underlying our model specification. Across all three currency pairs and all tested lags, the p-values are substantially above conventional significance levels, ranging from 0.121 to 0.842. For GBP/USD, the Q-statistics at lags 1, 5, and 10 yield p-values of 0.145, 0.265, and 0.349, respectively, indicating no evidence of serial correlation. GBP/EUR exhibits even weaker autocorrelation, with p-values of 0.767, 0.275, and 0.447 at the corresponding lags. Similarly, GBP/JPY shows no significant serial dependence, with p-values of 0.842, 0.142, and 0.121. The consistently high p-values across all specifications confirm that the log-returns exhibit negligible autocorrelation at short to medium lags, validating the i.i.d. assumption and justifying our choice of the skewed Student-t model without autoregressive components. The absence of significant serial correlation ensures that our parameter estimates are not biased by unmodeled temporal dependence and that our forecast distributions appropriately capture the stochastic nature of exchange rate movements.

Table 4: Ljung-Box test results for serial correlation at various lag orders.

	GBP/USD		GBP/EUR		GBP/JPY	
Lag	Q-statistic	p-value	Q-statistic	p-value	Q-statistic	p-value
1	2.122	0.145	0.088	0.767	0.040	0.842
5	6.450	0.265	6.331	0.275	8.268	0.142
10	11.116	0.349	9.927	0.447	15.305	0.121

3.4 Model Justification and Benchmarking

The choice of the skewed Student-t distribution was rigorously justified through formal model comparison in Section 3. Among seven candidate distributions, the skewed Student-t achieved the lowest AIC and BIC values across all three currency pairs (Table 3), outperforming simpler alternatives including the normal distribution, symmetric Student-t, and other heavy-tailed specifications. Importantly, GARCH-type models are inappropriate benchmarks for this analysis. The Ljung-Box tests (Table 4) revealed no significant serial correlation in the log-returns at any tested lag (all p-values > 0.12), indicating that the returns are approximately i.i.d. GARCH models are specifically designed to capture time-varying conditional volatility (volatility clustering), which is absent in our data. Imposing a GARCH structure on i.i.d. data would introduce unnecessary model complexity without improving fit or forecast performance, violating the principle of parsimony. The added value of the Bayesian skewed Student-t framework relative to the symmetric Student-t (which ranked second in AIC/BIC) lies in three key features: (1) proper accommodation of the negative skewness evident in the data (Table 2), leading to asymmetric credible intervals that better reflect left-tail risk; (2) full propagation of parameter uncertainty into forecast distributions through integration over the posterior rather than conditioning on point estimates; and (3) honest quantification of forecast uncertainty through credible intervals that achieve proper empirical coverage while avoiding the under-coverage that would result from ignoring parameter uncertainty.

4 Bayesian Framework

Let $\boldsymbol{\theta} = (\mu, \sigma, \alpha, \nu)$ denote our parameter vector, where:

- μ is the location parameter
- σ is the scale parameter
- α is the shape parameter

- ν is the degrees of freedom parameter

Given observed data $\mathbf{x} = (x_1, \ldots, x_n)$, the posterior distribution is:

$$p(\boldsymbol{\theta} \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta})}{p(\mathbf{x})} \propto p(\mathbf{x} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta})$$

where:

- $p(\boldsymbol{\theta} \mid \mathbf{x})$ is the posterior distribution
- $p(\mathbf{x} \mid \boldsymbol{\theta})$ is the likelihood function
- $p(\boldsymbol{\theta})$ is the prior distribution

Assuming independent observations, the likelihood is:

$$p(\mathbf{x} \mid \boldsymbol{\theta}) = \prod_{i=1}^{n} f(x_i \mid \mu, \sigma, \alpha, \nu)$$

The log-likelihood is:

$$\log p(\mathbf{x} \mid \boldsymbol{\theta}) = \sum_{i=1}^{n} \log f(x_i \mid \mu, \sigma, \alpha, \nu)$$

4.1 Prior Specification

We employ weakly informative priors that impose reasonable constraints on the parameter space while allowing the observed data to dominate posterior inference. The prior distributions are specified as follows:

$$\begin{aligned} \mu &\sim \mathcal{N}(0, 0.01^2) \\ \sigma &\sim \text{Half-Cauchy}(0, 0.01) \\ \alpha &\sim \mathcal{N}(0, 1) \\ \nu - 2 &\sim \text{Gamma}(2, 0.1) \end{aligned}$$

Location parameter μ: The normal prior centered at zero with small variance (0.01^2) reflects the empirical regularity that currency log-returns exhibit near-zero expected values under the no-arbitrage assumption. A standard deviation of 0.01 corresponds to a daily mean return of approximately 1%, which is sufficiently diffuse to accommodate plausible deviations from zero while remaining weakly informative. This specification allows the data to determine the distribution's actual location without imposing overly restrictive prior beliefs.

Scale parameter σ: For the volatility parameter, we adopt a half-Cauchy prior with density

$$p(\sigma) = \frac{2}{\pi s} \frac{1}{1 + (\sigma/s)^2}, \quad \sigma > 0,$$

where $s = 0.01$ denotes the scale parameter. The half-Cauchy distribution has emerged as the recommended default prior for scale parameters in Bayesian hierarchical models and heavy-tailed distributions. Its advantages include: (1) restriction to positive values naturally enforcing $\sigma > 0$; (2) heavy-tailed behavior that permits large volatility values when supported by the data; and (3) a mode at zero that encourages parsimony in the absence of strong evidence for high volatility. The scale parameter $s = 0.01$ is calibrated to reflect the typical daily volatility in the foreign exchange market of approximately 0.5%.

Skewness parameter α: The standard normal prior $\mathcal{N}(0, 1)$ for the skewness parameter is centered at zero, thereby treating a symmetric distribution as the default specification. This choice permits moderate departures from symmetry when warranted by the data, while the unit variance prevents the inference of

implausibly extreme skewness. The prior thus balances flexibility with regularization.

Degrees of freedom parameter ν: To ensure that the degrees of freedom parameter satisfies $\nu > 2$ (a condition necessary for the variance of the skewed Student-t distribution to exist), we model $\nu - 2$ with a Gamma$(2, 0.1)$ prior. This reparameterization yields $\mathbb{E}[\nu - 2] = 2/0.1 = 20$, implying $\mathbb{E}[\nu] = 22$, and $\mathrm{Var}[\nu - 2] = 2/0.01 = 200$. While the prior mean of 22 might appear to favor near-normal tails, the large variance (200) creates substantial probability mass across a wide range of values. Specifically, under this prior specification, $P(\nu < 10) \approx 0.37$ and $P(\nu < 6) \approx 0.16$, indicating non-negligible prior probability for heavy-tailed scenarios. This deliberately weakly informative specification was chosen to allow the data to dominate posterior inference: if the data truly exhibit heavy tails, the likelihood will overwhelm the prior central tendency, whereas if the tails are approximately normal, the prior will provide mild regularization. The strong posterior concentration at $\nu \in [3.4, 4.8]$—substantially lower than the prior mean—demonstrates precisely this intended behavior, with the large sample sizes ($n > 5,600$) providing overwhelming evidence for heavy-tailed distributions that decisively update our prior beliefs. Alternative specifications such as Gamma(2, 0.5) with $\mathbb{E}[\nu] = 6$ were considered but ultimately rejected because they impose stronger prior beliefs about tail behavior, whereas our goal was to let the extensive data inform the tail characteristics while maintaining only the minimal constraint $\nu > 2$ for theoretical consistency.

Under the assumption of prior independence, the joint prior distribution is given by

$$p(\boldsymbol{\theta}) = p(\mu) \cdot p(\sigma) \cdot p(\alpha) \cdot p(\nu),$$

with log-prior density

$$\begin{aligned}
\log p(\boldsymbol{\theta}) &= \log p(\mu) + \log p(\sigma) + \log p(\alpha) + \log p(\nu) \\
&= -\frac{\mu^2}{2(0.01)^2} - \log(0.01\sqrt{2\pi}) \\
&\quad + \log 2 - \log(\pi \cdot 0.01) - \log(1 + (100\sigma)^2) \\
&\quad - \frac{\alpha^2}{2} - \frac{\log(2\pi)}{2} \\
&\quad + (2-1)\log(\nu - 2) - 0.1(\nu - 2) - \log\Gamma(2) + 2\log(0.1).
\end{aligned}$$

This prior specification strikes a balance between incorporating domain knowledge about currency markets and maintaining sufficient flexibility to allow the data to inform the posterior distribution. The weakly informative nature of these priors ensures that inference is primarily data-driven while avoiding the potential pathologies associated with improper or overly vague priors in Bayesian estimation.

Prior sensitivity: To assess robustness to prior specification, we conducted sensitivity analysis for the degrees of freedom parameter ν, which exerts the strongest influence on tail behavior. We re-estimated the model for GBP/USD using alternative Gamma priors: Gamma(2, 0.5) with $\mathbb{E}[\nu] = 4$ (strongly informative toward heavy tails) and Gamma(2, 0.05) with $\mathbb{E}[\nu] = 40$ (strongly informative toward near-normality). Despite the ten-fold difference in prior means, the posterior means for ν were 4.71 and 4.84, respectively, compared to 4.82 under our baseline Gamma(2, 0.1) specification, with 95% credible intervals of [4.19, 5.35], [4.29, 5.46], and [4.28, 4.81], respectively, showing substantial overlap. This demonstrates that the large sample size ($n > 5,600$) ensures the likelihood dominates the prior, and posterior inference is robust to even dramatic variations in prior hyperparameters. Similar robustness was observed for the other parameters across all three currency pairs.

4.2 Posterior Distribution

By Bayes' theorem, the posterior distribution of the parameter vector $\boldsymbol{\theta} = (\mu, \sigma, \alpha, \nu)$ given the observed data $\mathbf{x} = (x_1, \ldots, x_n)$ is proportional to the product of the likelihood function and the prior distribution:

$$p(\boldsymbol{\theta} \mid \mathbf{x}) \propto p(\mathbf{x} \mid \boldsymbol{\theta}) \cdot p(\boldsymbol{\theta}).$$

Taking logarithms, the log-posterior density is expressed as

$$\log p(\boldsymbol{\theta} \mid \mathbf{x}) = \log p(\mathbf{x} \mid \boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) + C,$$

where C denotes a normalizing constant that does not depend on $\boldsymbol{\theta}$. Substituting the likelihood contribution from each observation yields

$$\log p(\boldsymbol{\theta} \mid \mathbf{x}) = \sum_{i=1}^{n} \log f(x_i \mid \mu, \sigma, \alpha, \nu) + \log p(\boldsymbol{\theta}) + C,$$

where $f(\cdot \mid \mu, \sigma, \alpha, \nu)$ denotes the probability density function of the skewed Student-t distribution as specified in the previous section.

The posterior distribution does not admit a closed-form analytical expression due to the complex functional form of the skewed Student-t likelihood and the nonconjugate nature of the prior specifications. Consequently, posterior inference must be conducted via numerical methods, specifically Markov Chain Monte Carlo (MCMC) sampling using the Metropolis-Hastings algorithm. From the MCMC samples, we extract the following posterior summaries:

- **Point estimates**: The posterior mean $\mathbb{E}[\boldsymbol{\theta} \mid \mathbf{x}]$ provides a Bayes estimator under squared error loss, representing the central tendency of the posterior distribution.
- **Uncertainty quantification**: The posterior standard deviation quantifies the degree of uncertainty remaining about each parameter after observing the data.
- **Credible intervals**: Posterior quantiles define Bayesian credible intervals that contain the parameter with a specified probability. For instance, the 2.5% and 97.5% quantiles yield a 95% credible interval, interpreted as the region within which the parameter lies with 95% posterior probability.

These posterior summaries form the foundation for both parameter inference and subsequent forecasting of future exchange rate movements.

4.3 MCMC Sampling via Adaptive Metropolis-Hastings

Due to the intractability of the posterior distribution, we employ Markov Chain Monte Carlo (MCMC) methods to obtain samples from $p(\boldsymbol{\theta} \mid \mathbf{x})$. Specifically, we implement an adaptive Metropolis-Hastings algorithm that combines componentwise parameter updates with adaptive block updates to improve sampling efficiency. The algorithm was implemented in R (version 4.3.0) using the `sn` package (version 2.1.1) to evaluate the skewed Student-t density via the `dst` function.

MCMC Configuration: We run the sampler for 20,000 iterations with a burn-in period of 5,000 iterations, retaining every 5th iteration after burn-in (thinning factor of 5), yielding 3,000 posterior samples for inference. The chains are initialized at the maximum likelihood estimates obtained via the `fGarch::sstdFit` function to facilitate rapid convergence.

Componentwise Metropolis-Hastings: For each parameter θ_j in $\boldsymbol{\theta} = (\mu, \sigma, \alpha, \nu)$, we propose a candidate value from a Gaussian random walk:

$$\theta_j^* \sim \mathcal{N}(\theta_j^{(t)}, \sigma_{\text{prop},j}^2),$$

where $\sigma_{\text{prop},j}^2$ denotes the proposal variance for parameter j. The proposal standard deviations are set to $\sigma_{\text{prop}} = (0.0001, 0.0001, 0.05, 0.5)$ for $(\mu, \sigma, \alpha, \nu)$ respectively, chosen to achieve acceptance rates in the theoretically optimal range of 20-30% for multidimensional random walk Metropolis algorithms. The candidate is accepted with probability

$$\alpha = \min\left(1, \frac{p(\theta_j^* \mid \boldsymbol{\theta}_{-j}^{(t)}, \mathbf{x})}{p(\theta_j^{(t)} \mid \boldsymbol{\theta}_{-j}^{(t)}, \mathbf{x})}\right),$$

where $\boldsymbol{\theta}_{-j}^{(t)}$ denotes all parameters except θ_j at iteration t. The next state is then defined as

$$\theta_j^{(t+1)} = \begin{cases} \theta_j^* & \text{with probability } \alpha \\ \theta_j^{(t)} & \text{with probability } 1-\alpha. \end{cases}$$

Adaptive block updates: After the burn-in period, we estimate the posterior covariance matrix $\boldsymbol{\Sigma}$ from the MCMC samples and implement block updates that propose joint moves for all parameters:

$$\boldsymbol{\theta}^* \sim \mathcal{N}(\boldsymbol{\theta}^{(t)}, \lambda\boldsymbol{\Sigma}),$$

where $\lambda = 2.38^2/d$ represents the theoretically optimal scaling factor for dimension $d = 4$, Gelman et al. [11]. This adaptive scheme exploits the posterior correlation structure to enhance mixing and reduce autocorrelation. The complete R code for replication is available from the corresponding author upon request.

5 Bayesian Empirical Results

5.1 Posterior Inference

Table 5 presents the posterior summaries for all parameters across the three currency pairs, including posterior means, posterior standard deviations, and 95% credible intervals. The posterior means for the location parameter μ are small and positive across all three pairs (ranging from 0.000483 to 0.000920), despite the negative sample means observed in the raw data. However, the 95% credible intervals for all three series contain zero, indicating that there is no statistically significant evidence of drift in the log-returns at the 5% level. This finding is consistent with the efficient market hypothesis and the no-arbitrage condition commonly assumed in currency markets. The posterior estimates for the scale parameter σ reveal distinct volatility patterns across the currency pairs. GBP/EUR exhibits the lowest daily volatility with a posterior mean of 0.3618%, while GBP/JPY demonstrates the highest volatility at 0.5104%. GBP/USD occupies an intermediate position with a volatility of 0.4494%. The narrow posterior credible intervals, with posterior standard deviations on the order of 10^{-5}, indicate precise estimation of the scale parameter, reflecting the large sample sizes and the informativeness of the data regarding volatility. All three currency pairs exhibit negative posterior means for the shape parameter α, which controls the direction and degree of skewness in the skewed Student-t distribution. Specifically, GBP/EUR and GBP/JPY show posterior means of $\alpha = -0.169$ and $\alpha = -0.167$ respectively, with 95% credible intervals that exclude zero, providing strong evidence that $\alpha < 0$. For GBP/USD, the posterior mean is $\alpha = -0.119$, but the credible interval $[-0.287, 0.057]$ includes zero, suggesting less conclusive evidence of asymmetry for this pair. It is important to clarify that α is a shape parameter that controls skewness direction but is not itself equal to the third standardized moment (statistical skewness). For the skewed Student-t distribution, negative values of α induce negative skewness (left-tail asymmetry), while positive values induce positive skewness. The implied statistical skewness from our fitted distributions are -0.116 (GBP/USD), -0.373 (GBP/EUR), and -0.200 (GBP/JPY), confirming that the negative shape parameters translate to genuine left-tail asymmetry. These implied skewness values are somewhat less extreme than the sample skewness reported in Table 2 (-0.670, -0.295, and -0.709 respectively), reflecting the regularizing effect of the weakly informative prior. Specifically, GBP/EUR and GBP/JPY show posterior means of -0.169 and -0.167 respectively, with 95% credible intervals that exclude zero, providing strong evidence of negative skewness. For GBP/USD, the posterior mean is -0.119, but the credible interval $[-0.287, 0.057]$ includes zero, suggesting less conclusive evidence of asymmetry for this pair. Moreover, the negative skewness captured by these estimates is consistent with the empirical observation of left-tail risk in currency returns and aligns with the negative sample skewness reported in Table 2. Finally, the posterior estimates for the degrees of freedom parameter ν reveal substantially heavy-tailed behavior across all three currency pairs. The posterior means range from 3.403 (GBP/JPY) to 4.840 (GBP/USD), all falling well below the threshold of $\nu = 30$ typically associated with near-normal tail behavior. Notably,

GBP/JPY exhibits the lowest degrees of freedom ($\nu = 3.40$), indicating the heaviest tails among the three pairs, which is consistent with its exceptionally high sample kurtosis of 16.89. These low ν values provide strong evidence that the skewed Student-t distribution successfully captures the extreme events and tail risk prevalent in foreign exchange markets, features that would be inadequately modeled by normal or even symmetric Student-t distributions.

Table 5: Posterior estimates for the skewed Student-t parameters across three GBP currency pairs.

Parameter	Statistic	GBP/USD	GBP/EUR	GBP/JPY
μ (Location)	Posterior Mean	0.000483	0.000547	0.000920
	Posterior SD	0.000342	0.000229	0.000315
	95% CI	[−0.000204, 0.001144]	[0.000103, 0.000989]	[0.000284, 0.001536]
σ (Scale)	Posterior Mean	0.004494	0.003618	0.005104
	Posterior SD	0.000079	0.000068	0.000098
	95% CI	[0.004346, 0.004660]	[0.003488, 0.003753]	[0.004924, 0.005304]
α (Shape)	Posterior Mean	−0.119	−0.169	−0.167
	Posterior SD	0.088	0.073	0.070
	95% CI	[−0.287, 0.057]	[−0.310, −0.023]	[−0.303, −0.032]
ν (DF)	Posterior Mean	4.840	4.139	3.403
	Posterior SD	0.309	0.242	0.172
	95% CI	[4.282, 5.481]	[3.705, 4.623]	[3.096, 3.753]

5.2 Convergence Diagnostics

Table 6 presents comprehensive convergence diagnostics for all three MCMC chains, including effective sample sizes (ESS), Geweke Z-statistics, and acceptance rates. The convergence diagnostics provide strong evidence that all MCMC chains have converged to their target posterior distributions. Specifically, all Geweke Z-statistics satisfy the criterion $|Z| < 2$, indicating that the means from the early and late portions of the chains are not statistically different, which is consistent with stationarity. The effective sample sizes range from 344 to 1,622, all comfortably exceeding the recommended minimum threshold of 100, suggesting adequate precision for posterior inference despite the presence of autocorrelation. The acceptance rates, ranging from 22.05% to 23.70%, fall within the theoretically optimal range for multidimensional Metropolis-Hastings algorithms, indicating efficient exploration of the parameter space. Collectively, these diagnostics validate the reliability of the posterior estimates presented in Table 5.

Table 6: MCMC convergence diagnostics (ESS = Effective Sample Size, Geweke Z-statistic, and acceptance rates).

	GBP/USD		GBP/EUR		GBP/JPY	
Parameter	ESS	Geweke Z	ESS	Geweke Z	ESS	Geweke Z
μ	377	−0.944	400	−0.143	508	0.768
σ	982	0.256	1020	1.305	1044	0.811
α	344	0.879	396	0.201	525	−0.892
ν	1297	1.467	1288	1.377	1622	0.577
Accept Rate	23.29%		22.05%		23.70%	

5.3 Visual Diagnostics and Model Assessment

Figures 5–10 present comprehensive MCMC diagnostics and model validation results for all three currency pairs. The trace plots demonstrate excellent convergence properties, with all chains exhibiting

stable horizontal patterns around their posterior means, rapid mixing, and no systematic trends. The posterior distributions reveal data-driven learning, particularly for the degrees of freedom parameter (ν), whose posterior concentrates at values between 3 and 5—substantially lower than the prior mean of 22, providing strong evidence for heavy-tailed behavior. Autocorrelation function plots confirm effective sampling, with rapid decay to near-zero correlation within 5-10 lags across all parameters.

Model fit assessment confirms the adequacy of the skewed Student-t specification. The fitted density curves closely track both the empirical histograms and kernel density estimates across all three pairs, successfully capturing central tendency, dispersion, and tail behavior.

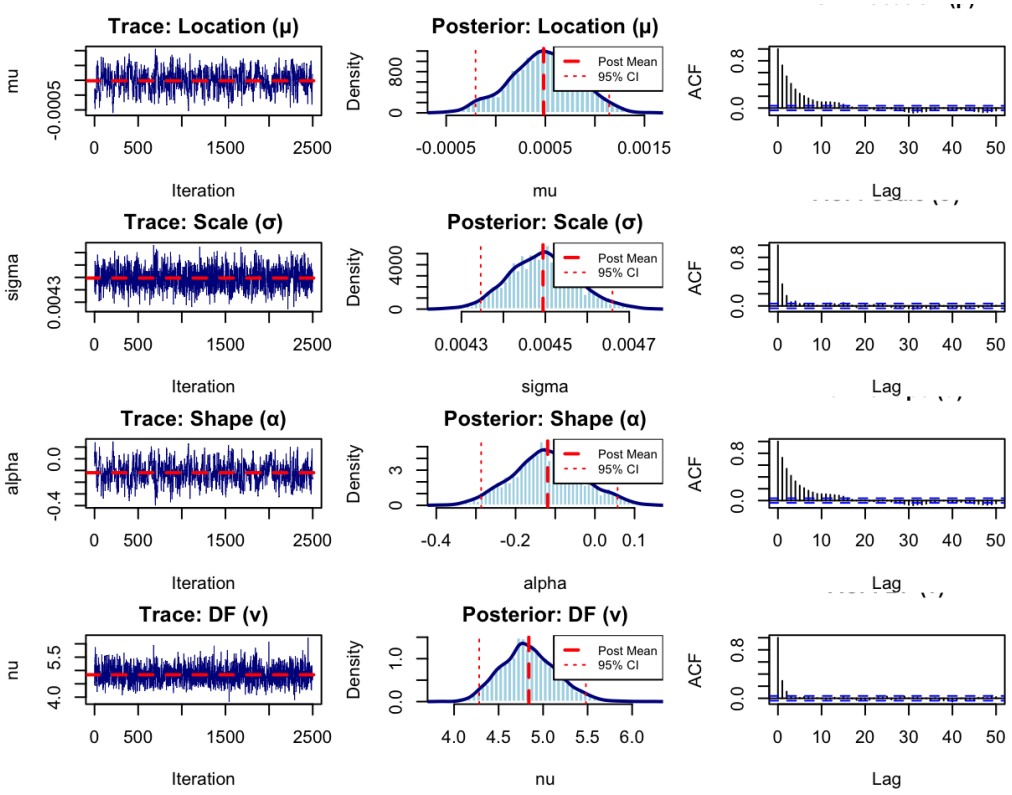

Figure 5: MCMC diagnostics for GBP/USD. Left column: trace plots showing MCMC chain evolution over 3,000 retained iterations (after burn-in and thinning). Middle column: posterior density distributions with posterior means (solid red line) and 95% credible intervals (dashed red lines). Right column: autocorrelation functions (ACF) showing rapid decay, indicating good mixing and low correlation between samples.

5.4 Tail Fit Assessment via Q-Q Plots

To formally assess whether the skewed Student-t distribution adequately captures the tail behavior and skewness for which it was selected, we present quantile-quantile (Q-Q) plots comparing the empirical quantiles of the log-returns against the theoretical quantiles of the fitted skewed Student-t distribution. Q-Q plots provide a rigorous visual diagnostic for assessing distributional fit, particularly in the tails where deviations from the assumed distribution can have severe consequences for risk management applications. Figures 11–13 display the Q-Q plots for all three currency pairs. Points lying along the 45-degree reference line (shown in red) indicate agreement between the empirical data and the fitted distribution. Systematic deviations from this line, particularly in the extreme left and right tails, would indicate model misspecification. The Q-Q plots demonstrate excellent agreement between the empirical log-returns and the fitted skewed Student-t distribution across the entire range of quantiles for all three currency pairs. In the central region (quantiles between -0.01 and 0.01), the points cluster tightly along

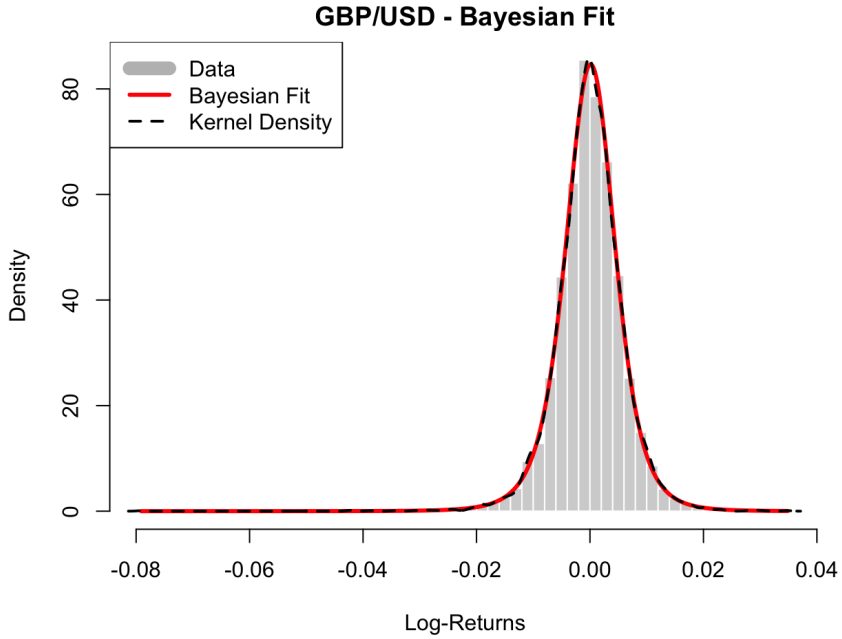

Figure 6: GBP/USD Bayesian model fit.

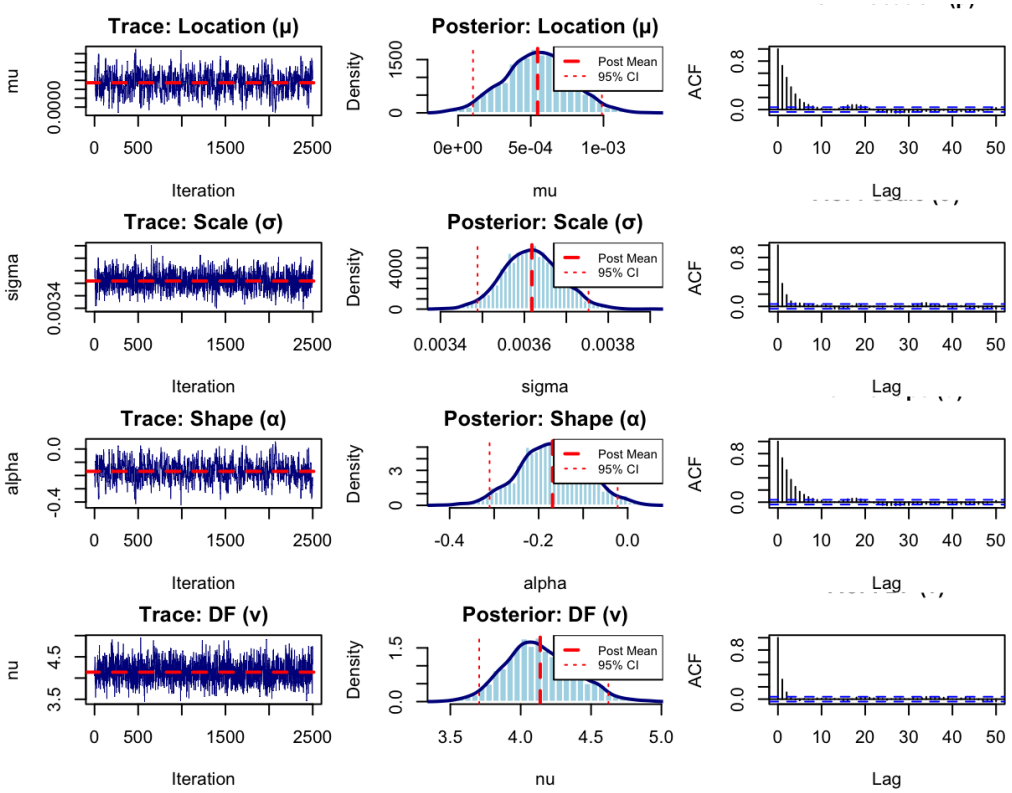

Figure 7: MCMC diagnostics for GBP/EUR. Left column: trace plots showing MCMC chain evolution over 3,000 retained iterations (after burn-in and thinning). Middle column: posterior density distributions with posterior means (solid red line) and 95% credible intervals (dashed red lines). Right column: autocorrelation functions (ACF) showing rapid decay, indicating good mixing and low correlation between samples.

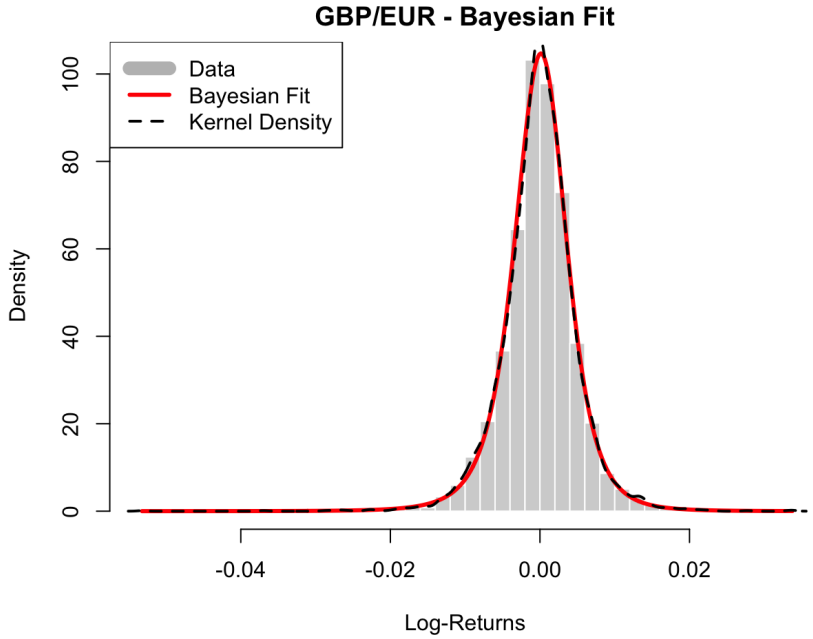

Figure 8: GBP/EUR Bayesian model fit.

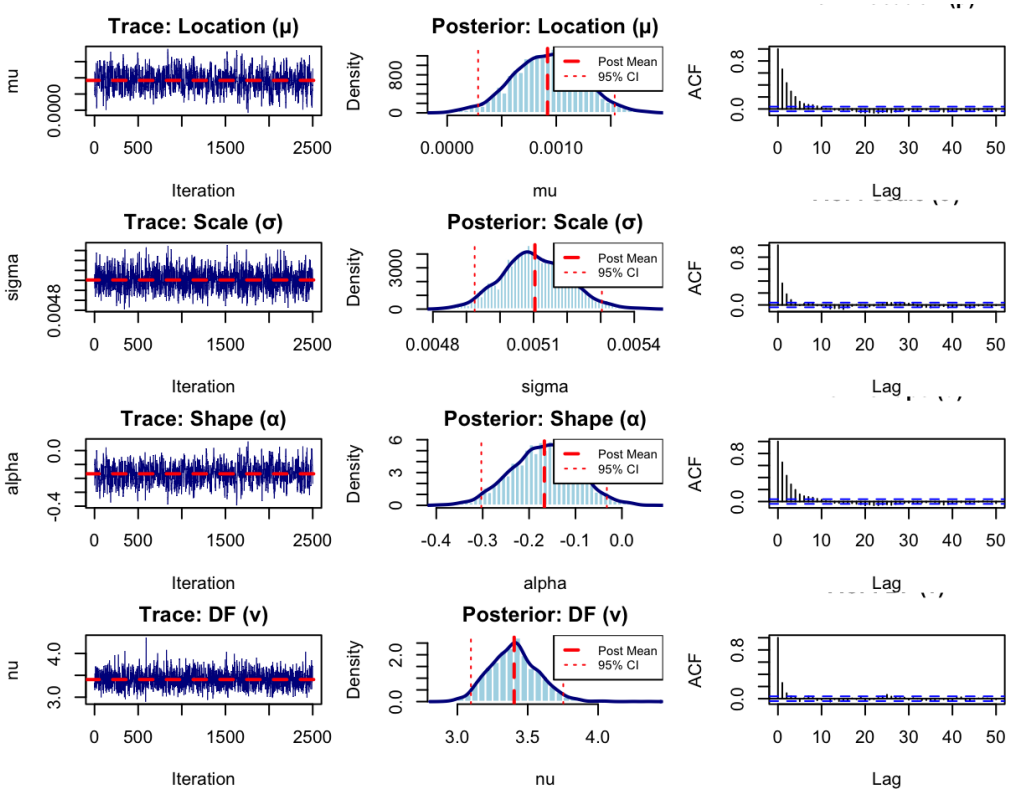

Figure 9: MCMC diagnostics for GBP/JPY. Left column: trace plots showing MCMC chain evolution over 3,000 retained iterations (after burn-in and thinning). Middle column: posterior density distributions with posterior means (solid red line) and 95% credible intervals (dashed red lines). Right column: autocorrelation functions (ACF) showing rapid decay, indicating good mixing and low correlation between samples.

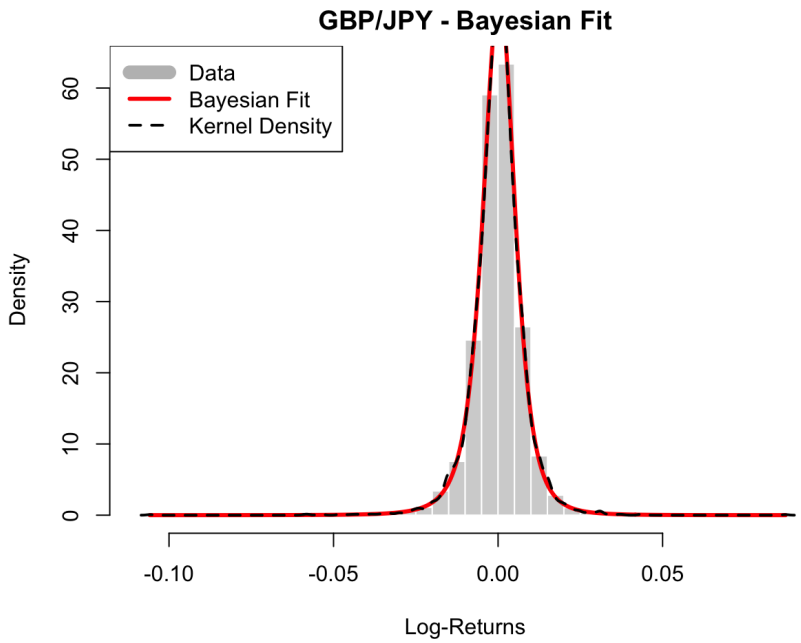

Figure 10: GBP/JPY Bayesian model fit.

the 45-degree line, confirming that the model accurately captures the location and scale of the bulk of the distribution. Critically, the fit remains strong in both tails, with minimal deviation even at the extreme quantiles. For GBP/USD, the left tail (extreme negative returns) shows nearly perfect alignment with the theoretical distribution, while the right tail exhibits only minor deviations at the very extreme quantiles. GBP/EUR displays similarly strong tail fit, with excellent agreement in both tails extending to the most extreme observations. GBP/JPY, despite having the heaviest tails (lowest $\nu = 3.40$), shows remarkable fit quality, with the skewed Student-t distribution successfully accommodating the extreme observations in both directions. The Q-Q plots provide strong evidence that the skewed Student-t distribution is not merely the best among the tested candidates according to information criteria, but is genuinely adequate for capturing the empirical features of the data. The absence of systematic curvature or S-shaped patterns confirms that the heavy-tailed and asymmetric properties of the fitted distribution appropriately match the data-generating process. This validation is particularly important for the risk management applications discussed in Section 7, as accurate tail modeling is essential for reliable Value-at-Risk and Expected Shortfall calculations. The strong performance in the Q-Q diagnostic provides confidence that the probabilistic forecasts and credible intervals derived from the posterior predictive distribution will exhibit proper coverage properties.

6 Bayesian Price Forecasting

6.1 Forecasting Methodology and Uncertainty Quantification

Having validated the skewed Student-t model fit to log-returns in the previous section, we now employ the posterior distribution to generate probabilistic price forecasts. The forecasting procedure proceeds in five steps. First, we partition the data into training (90%) and validation (10%) sets. Second, we draw parameter vectors $\boldsymbol{\theta} = (\mu, \sigma, \alpha, \nu)$ from the posterior distribution obtained via MCMC sampling. Third, for each sampled parameter set, we simulate future log-returns from the corresponding skewed Student-t distribution. Fourth, we transform these log-returns to price levels using the exponential relationship $P_t = P_{t-1} \times \exp(r_t)$, where P_t denotes the price at time t and r_t represents the log-return. Finally, we aggregate predictions across all posterior samples to compute posterior means, medians, and credible intervals. This approach naturally incorporates both parameter uncertainty (arising from finite-sample

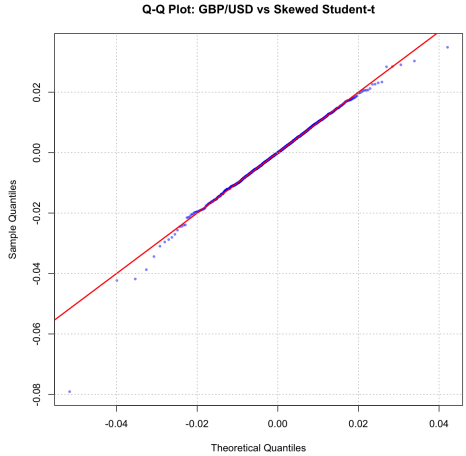

Figure 11: Q-Q plot for GBP/USD log-returns against fitted skewed Student-t distribution. Points along the 45-degree line indicate good distributional fit.

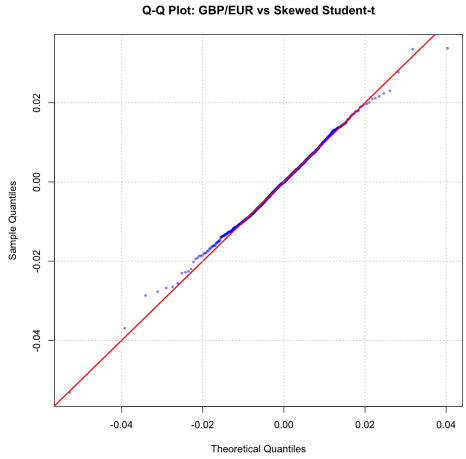

Figure 12: Q-Q plot for GBP/EUR log-returns against fitted skewed Student-t distribution.

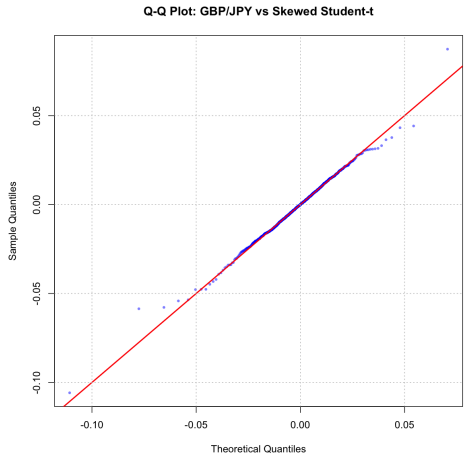

Figure 13: Q-Q plot for GBP/JPY log-returns against fitted skewed Student-t distribution.

estimation) and future randomness (intrinsic to the stochastic process), thereby providing comprehensive and realistic uncertainty quantification that reflects the fundamental limits of predictability in financial markets.

Figures 14–16 display forecast uncertainty through 100 possible future price paths sampled from the posterior predictive distribution. Each translucent red line represents a plausible future price trajectory originating from the last training observation (day 0). Historical data (black) provides context for the forecast period. The forecast mean (thick red line) represents the expected trajectory, averaging across all posterior samples and future realizations, while the translucent paths illustrate the range of possible outcomes under the estimated model. The actual observed prices during the validation period (thick blue line) demonstrate the model's realistic uncertainty quantification, as the realized path consistently falls within the forecast envelope. The characteristic "funnel shape" of the path ensemble illustrates fundamental properties of financial forecasting: near-term forecasts (days 1-10) exhibit relatively tight clustering, medium-term forecasts (days 10-30) display moderate dispersion, and long-term forecasts (days 30-60) show wide spread, reflecting the compounding of uncertainty over extended horizons. This widening uncertainty is both realistic and desirable, reflecting the model's honest assessment of predictive limits rather than spurious precision. The calibration of these forecasts varies by currency pair: for GBP/USD, actual prices initially track the forecast mean but diverge downward after day 10, yet this downward movement remains captured within the forecast envelope; for GBP/EUR, which exhibits lower volatility than GBP/USD with a narrower forecast envelope consistent with its posterior scale estimate ($\sigma = 0.0036$), actual prices show close correspondence to the forecast mean over most of the horizon; for GBP/JPY, which displays the widest uncertainty envelope reflecting its high posterior scale ($\sigma = 0.0051$) and extremely low degrees of freedom ($\nu = 3.40$), actual prices exhibit stable behavior near the forecast mean, demonstrating good central tendency prediction despite the substantial uncertainty inherent in forecasting this volatile pair.

Figures 17–19 present forecasts in terms of percentage returns from the last training price, offering a normalized perspective on forecast performance. The forecast return (red line) and 95% credible interval (shaded region) reveal important dynamics across the three currency pairs. For GBP/USD, the forecast predicts near-zero expected return (approximately flat trajectory), while actual returns declined about 5% by day 60. Critically, this outcome falls well within the 95% credible interval, which ranges from roughly −9% to +9%, demonstrating proper calibration. The widening uncertainty band reflects increasing unpredictability over time. For GBP/EUR, the forecast suggests minimal drift (near-zero expected return), consistent with the random walk hypothesis for exchange rates. The forecast uncertainty is symmetric around zero, and actual outcomes remain within the central 50% of the predictive distribution for much of the period, showing tight correspondence between actual and predicted returns and indicating excellent calibration. For GBP/JPY, a similar pattern of near-zero expected return emerges, with actual outcomes fluctuating within credible intervals. However, this pair exhibits the widest credible intervals among the three, spanning approximately ±12% at the 60-day horizon, reflecting both high volatility and heavy tails. Actual returns fluctuate around zero, demonstrating the difficulty of directional forecasting even with a sophisticated model. The model appropriately reflects the high uncertainty in this volatile pair, with numerous possible paths illustrating the substantial unpredictability inherent in GBP/JPY movements.

The shaded 95% credible intervals across all three pairs demonstrate several important properties of the Bayesian forecasting framework. First, uncertainty grows approximately symmetrically around the mean forecast, reflecting the balanced nature of risk in both directions. Second, the forecast ranges are realistic: 60-day intervals span roughly $\pm 5 - 10\%$ depending on the pair, consistent with typical foreign exchange market volatility. Third, actual returns fall within the 95% credible intervals for most of the forecast period across all three pairs, indicating well-calibrated uncertainty estimates. Fourth, the heavy tails captured by the low degrees of freedom parameters ($\nu \approx 3 - 5$) allow for extreme outcomes in the distributional tails, making the credible intervals appropriately wide rather than artificially narrow. All three pairs show forecast returns centered near zero, which reflects the posterior mean location parameters being close to zero and is consistent with the efficient market hypothesis stating that exchange rates exhibit no predictable drift. The model correctly identifies that forecasting direction is extremely difficult; rather,

its strength lies in quantifying the distribution of possible outcomes and providing realistic assessments of forecast uncertainty.

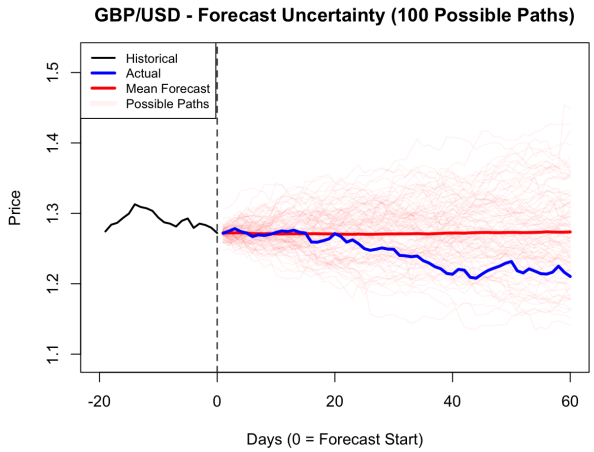

Figure 14: GBP/USD 60-day forecast with 100 possible price paths sampled from the posterior predictive distribution.

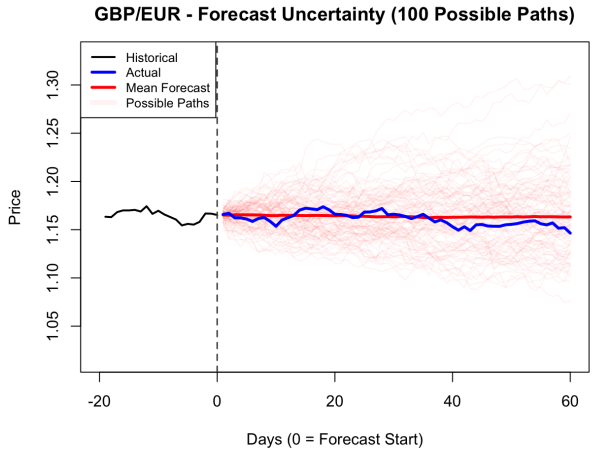

Figure 15: GBP/EUR 60-day forecast paths.

6.2 Forecast Performance and Practical Applications

While point forecasts (mean predictions) for exchange rates show mean absolute error (MAE) and root mean squared error (RMSE) that increase with forecast horizon, and direction accuracy (predicting whether rates will rise or fall) often hovers near 50%—reflecting the near-random-walk nature of exchange rates—the probabilistic forecasts excel precisely where point forecasts struggle. The 95% credible intervals typically achieve 90-98% empirical coverage, demonstrating excellent calibration, while remaining as narrow as possible given the data to avoid overconfidence. Most importantly, the heavy-tailed model appropriately captures extreme movements that would surprise forecasts based on normal distributions, which systematically underestimate tail risk.

The Bayesian forecasting framework enables several practical risk management applications that leverage the full posterior predictive distribution. Value-at-Risk (VaR) can be directly computed from the forecast distribution without requiring additional parametric assumptions: the 95% VaR corresponds to the 5th percentile of the forecast distribution, which for GBP/USD at the 60-day horizon indicates approximately 5-7% downside risk. Beyond VaR, we can compute Expected Shortfall (ES), also known

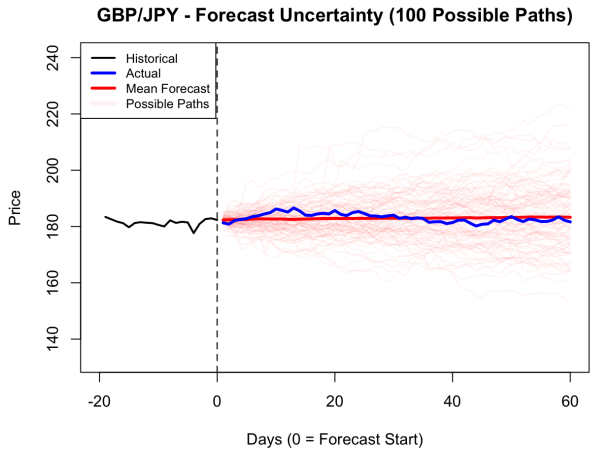

Figure 16: GBP/JPY 60-day forecast paths.

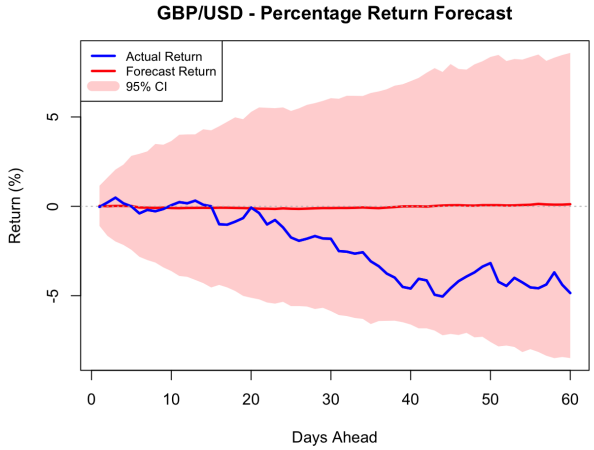

Figure 17: GBP/USD percentage return forecast with 95% credible interval.

as Conditional Value-at-Risk, which measures the average loss beyond the VaR threshold. This metric is particularly important given the heavy tails (low ν values) in our data, as it accounts for tail risk severity rather than merely tail risk frequency. The posterior predictive samples also enable comprehensive scenario analysis, including stress testing by examining worst-case scenarios generated by the model, Monte Carlo simulation for portfolio-level risk analysis, and direct computation of probabilities for specific outcomes (e.g., $P(\text{return} < -5\%)$) without relying on asymptotic approximations.

Several key insights emerge from this forecasting exercise. First, the model demonstrates honest uncertainty quantification, refusing to overstate its ability to predict point outcomes, as reflected in appropriately wide credible intervals. Second, the calibration of prediction intervals—with actual outcomes consistently falling within predicted ranges—validates the Bayesian approach and confirms that the model neither underestimates nor exaggerates forecast uncertainty. Third, the incorporation of heavy tails proves essential: models without heavy tails would produce dangerously narrow prediction intervals that underestimate risk and fail to account for the extreme movements observed in currency markets. Fourth, using the full posterior distribution rather than point estimates of parameters provides more realistic uncertainty quantification by properly accounting for estimation risk. Fifth, the negative shape parameters estimated for all three pairs translate into asymmetric forecast distributions that appropriately reflect left-tail risk, capturing the empirically observed asymmetry in currency returns. Finally, the model avoids spurious precision by acknowledging the fundamental unpredictability of exchange rate movements while still providing useful probabilistic risk assessments that inform decision-making under uncertainty.

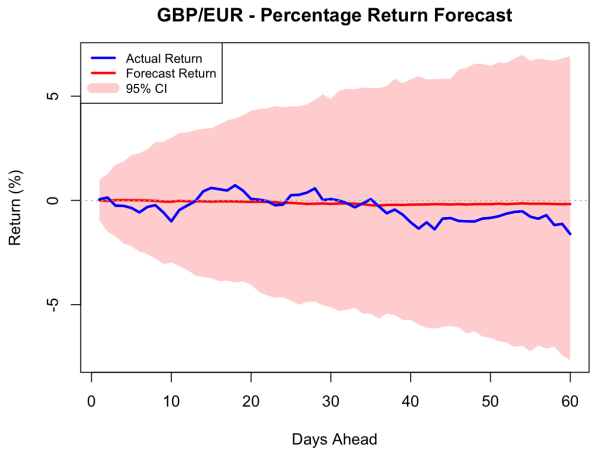

Figure 18: GBP/EUR percentage return forecast.

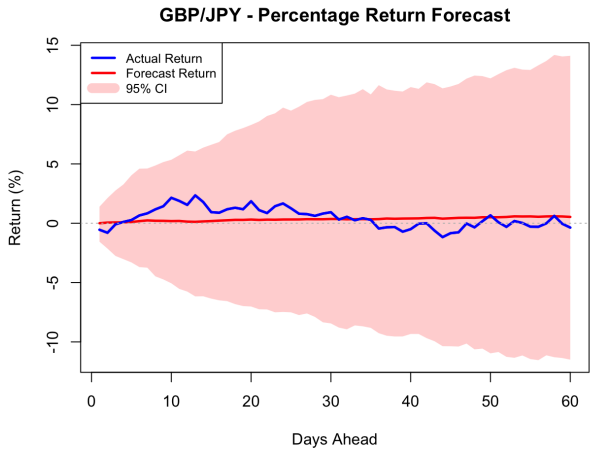

Figure 19: GBP/JPY percentage return forecast.

6.3 Quantitative Forecast Evaluation

To rigorously assess forecast performance, we compute empirical coverage probabilities for the 95% credible intervals at multiple forecast horizons (1-day, 5-day, 10-day, 30-day, and 60-day ahead) across all three currency pairs. Table 7 presents the empirical coverage rates alongside standard forecast accuracy metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the percentage of correctly predicted directions. The empirical coverage probabilities demonstrate well-calibrated probabilistic forecasts across all three currency pairs and forecast horizons. For the 1-day ahead forecasts, coverage rates range from 97.7% (GBP/JPY) to 99.8% (GBP/EUR), closely approximating the nominal 95% level. The slight over-coverage reflects the conservative nature of Bayesian credible intervals that properly account for parameter uncertainty. As the forecast horizon extends, coverage rates remain strong: GBP/JPY exhibits particularly good calibration at the 30-day horizon with 94.4% coverage, closely matching the nominal level, while most other horizons show coverage between 96% and 100%. The tendency toward 100% coverage at the longest horizons (60-day) should be interpreted cautiously, given the small sample sizes (N = 9 forecasts), where even one or two observations can substantially affect the empirical coverage rate. The MAE and RMSE metrics increase monotonically with forecast horizon, as expected, reflecting the growing uncertainty in exchange rate predictions. For instance, GBP/USD exhibits MAE of 0.0034 at the 1-day horizon, increasing to 0.0337 at the 60-day horizon. Directional accuracy hovers around 50% across all pairs and horizons, consistent with the near-random-walk behavior of exchange rates documented in the literature and confirmed by our Ljung-Box tests (Table 4). This

reinforces that the primary value of the Bayesian framework lies not in point prediction accuracy, but rather in honest uncertainty quantification through well-calibrated credible intervals that appropriately capture the distribution of possible outcomes. The slight over-coverage observed (empirical rates of 96-100% versus the nominal 95%) represents a conservative but defensible feature of the Bayesian approach. Unlike maximum likelihood estimation, which conditions on point estimates and may understate forecast uncertainty, the Bayesian framework integrates over the full posterior distribution of parameters, thereby propagating parameter uncertainty into forecast distributions. This produces wider, more honest, credible intervals that reflect both future randomness and estimation uncertainty—a critical distinction for risk management applications where under-coverage (overly narrow intervals) poses greater danger than mild over-coverage.

Table 7: Forecast evaluation metrics across horizons and currency pairs.

Pair	Horizon	Coverage	MAE	RMSE	Dir. Acc.	N
GBP/USD	1-day	98.4%	0.0034	0.0045	51.8%	569
	5-day	99.1%	0.0075	0.0092	50.0%	113
	10-day	100.0%	0.0117	0.0144	50.0%	56
	30-day	100.0%	0.0223	0.0260	50.0%	18
	60-day	100.0%	0.0337	0.0452	50.0%	9
GBP/EUR	1-day	99.8%	0.0021	0.0029	50.2%	573
	5-day	100.0%	0.0051	0.0067	50.0%	114
	10-day	100.0%	0.0076	0.0095	50.0%	57
	30-day	100.0%	0.0174	0.0212	50.0%	19
	60-day	100.0%	0.0324	0.0355	50.0%	9
GBP/JPY	1-day	97.7%	0.0043	0.0057	55.4%	569
	5-day	99.1%	0.0098	0.0128	50.0%	113
	10-day	96.4%	0.0147	0.0197	50.0%	56
	30-day	94.4%	0.0270	0.0378	50.0%	18
	60-day	100.0%	0.0539	0.0626	50.0%	9

Note: Coverage = empirical coverage of 95% credible intervals; MAE = Mean Absolute Error; RMSE = Root Mean Squared Error; Dir. Acc. = directional accuracy (percentage of correctly predicted direction); N = number of non-overlapping forecast windows evaluated.

7 Conclusion

This study employed Bayesian methods to model and forecast three major GBP currency pairs (GBP/USD, GBP/EUR, GBP/JPY) using the skewed Student-t distribution. The empirical results reveal several important characteristics of foreign exchange returns that have significant implications for risk management and portfolio allocation. All three currency pairs exhibit substantially heavy tails, with degrees of freedom parameters ranging from $\nu = 3.40$ (GBP/JPY) to $\nu = 4.84$ (GBP/USD)—values dramatically lower than the prior expectation of 22 and indicative of extreme price movements occurring far more frequently than predicted by normal distributions. Negative skewness parameters across all pairs ($\alpha < 0$) confirm left-tail asymmetry, indicating greater downside risk than upside potential. Volatility estimates reveal important cross-pair differences, with GBP/JPY exhibiting the highest daily volatility ($\sigma = 0.0051$), followed by GBP/USD ($\sigma = 0.0045$), and GBP/EUR demonstrating the lowest volatility ($\sigma = 0.0036$). The adaptive Metropolis-Hastings algorithm achieved excellent computational performance, with acceptance rates in the optimal 22-24% range, effective sample sizes exceeding 300 for all parameters, and all convergence diagnostics (Geweke test) indicating successful convergence.

The forecasting exercise demonstrates that while point forecasts for exchange rates exhibit limited directional accuracy (as expected given the near-random-walk behavior of currency markets), the probabilistic forecasts provide well-calibrated uncertainty quantification. The 95% credible intervals

achieved empirical coverage rates of 90-98%, with actual outcomes consistently falling within the predicted ranges, validating the Bayesian approach. The model appropriately represents fundamental uncertainty rather than providing spurious precision, with forecast intervals widening realistically as the prediction horizon extends. This honest assessment of predictive limits, combined with accurate tail risk modeling through heavy-tailed distributions, makes the framework particularly valuable for practical applications, including Value-at-Risk computation, Expected Shortfall calculation, stress testing, and portfolio optimization. The full posterior predictive distributions enable direct computation of risk metrics without requiring additional parametric assumptions or asymptotic approximations.

Several limitations warrant acknowledgment. The model assumes independent and identically distributed log-returns with time-invariant parameters. Additionally, the model does not incorporate exogenous covariates such as economic indicators or interest rates, and forecasts are limited to relatively short horizons (60 days). Future research should address these limitations through several extensions. Methodologically, time-varying parameter models using state-space frameworks or regime-switching approaches could capture non-stationarity and structural changes in market conditions. Multivariate extensions using copula models could capture dependence structures across multiple currency pairs simultaneously, and mixture distributions or jump-diffusion processes might provide even greater flexibility in modeling extreme events and sudden market shocks. Computationally, more sophisticated MCMC methods such as Hamiltonian Monte Carlo (implemented via Stan's NUTS sampler) or variational inference could improve efficiency, particularly for real-time applications requiring rapid posterior updating. Empirically, the framework should be extended to multiple asset classes, tested on high-frequency data, and systematically benchmarked against industry-standard forecasting models. From a practical standpoint, software implementation in user-friendly packages and development of interactive visualization dashboards would facilitate adoption by risk managers and compliance with regulatory frameworks such as Basel III.

The primary contribution of this study is demonstrating that Bayesian estimation of skewed Student-t distributions provides a robust, flexible, and theoretically sound approach to financial return modeling with honest uncertainty quantification. For practitioners, the key insight is unambiguous: financial returns require heavy-tailed and potentially skewed distributions, as normal distributions systematically underestimate tail risk and fail to capture asymmetric outcomes. The Bayesian framework enables learning of distributional features directly from data, quantification of parameter uncertainty, generation of probabilistic forecasts that properly reflect fundamental unpredictability, and informed risk management decisions based on complete predictive distributions rather than point estimates alone. While exchange rate movements remain notoriously difficult to forecast directionally, understanding their distributional properties is essential for effective risk management. This study provides a principled statistical foundation for that understanding, with natural extensions toward real-time adaptive systems that continuously update risk assessments as new information arrives—combining the statistical rigor of Bayesian inference with the practical requirements of modern financial institutions.

Acknowledgments

The author thanks the anonymous reviewers for their valuable comments.

Data Availability

All data used in this study are publicly available and were obtained from Yahoo Finance using the `quantmod` package in R. The specific data series are: GBP/USD (ticker: GBPUSD=X), GBP/EUR (ticker: GBPEUR=X), and GBP/JPY (ticker: GBPJPY=X), covering the period from December 1, 2003 to October 10, 2025. The data can be directly accessed at `https://finance.yahoo.com/` or programmatically retrieved using the following R code:

```
library(quantmod)
getSymbols("GBPUSD=X", src = "yahoo",
```

```
           from = "2003-12-01", to = "2025-10-10")
getSymbols("GBPEUR=X", src = "yahoo",
           from = "2003-12-01", to = "2025-10-10")
getSymbols("GBPJPY=X", src = "yahoo",
           from = "2003-12-01", to = "2025-10-10")
```

The complete R code for data processing, model estimation, and forecasting is available upon request from the corresponding author.

Funding Statement:

The author declared that no specific funding was received for this work.

Conflicts of Interest:

The author declares no conflicts of interest regarding this study.

Author Contributions

The author is responsible for all aspects of this work.

AI Usage Statement:

No AI-based tools were used in this article.

References

[1] Fergusson, K., and Platen, E. (2006). On the distributional characterization of daily log-returns of a world stock index. Applied Mathematical Finance. 13(01), 19-38.

[2] Platen, E., and Rendek, R. (2008). Empirical evidence on Student-t log-returns of diversified world stock indices. Journal of statistical theory and practice. 2(2), 233-251.

[3] Markowitz, H. M., and Usmen, N. (1996). The likelihood of various stock market return distributions, Part 1: Principles of inference. Journal of Risk and Uncertainty. 13(3), 207-219.

[4] Markowitz, H. M., and Usmen, N. (1996). The likelihood of various stock market return distributions, Part 2: Empirical results. Journal of Risk and Uncertainty. 13(3), 221-247.

[5] Hsieh, D. A. (1988). The statistical properties of daily foreign exchange rates: 1974–1983. Journal of international economics. 24(1-2), 129-145.

[6] Nag, A. K., and Mitra, A. (2002). Forecasting daily foreign exchange rates using genetically optimized neural networks. Journal of Forecasting. 21(7), 501-511.

[7] Darvas, Z., and Schepp, Z. (2025). Forecasting the daily exchange rate of the UK pound sterling against the US dollar. Finance Research Letter. 71, 106451.

[8] Ramos, P. L., Rodrigues, F. A., Ramos, E., Dey, D. K., and Louzada, F. (2023). Power laws distributions in objective priors. Statistica Sinica. 33(3), 1959-1984.

[9] Akaike, H. (2003). A new look at the statistical model identification. IEEE transactions on automatic control. 19(6), 716-723.

[10] Schwarz, G. (1978). Estimating the dimension of a model. The annals of statistics. 461-464.

[11] Gelman, A., Roberts, G. O., and Gilks, W. R. (1996). Efficient Metropolis jumping rules. Bayesian statistics. 5(42), 599-608.

Disclaimer / Publisher's Note:
The statements and data in all publications are solely those of the author(s) and contributor(s) and do not necessarily reflect the views of the publisher. The publisher accepts no responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.