Unilag Journal of Mathematics and Applications,
Volume 6, (2026), Pages 1--19.
ISSN: 2805 3966. URL: http://lagjma.edu.ng

A BAYESIAN HERMITE REGRESSION MODEL WITH CLASSES OF PRIOR DISTRIBUTIONS APPLIED TO RIDE-HAILING PLATFORM USAGE

TITILOYE KEHINDE ABAYOMI*

ABSTRACT. Econometric and applied statistical modelling frequently encounter nonlinear datasets characterised by heavy tails, skewness, and volatility clustering. Classical regression methods, including ordinary least squares, often perform poorly under such conditions, while conventional polynomial regression may be unstable in the presence of extreme observations. To address these limitations, this study develops a Bayesian Hermite Regression Model (BHRM) that integrates truncated Hermite polynomial expansions within a coherent Bayesian framework. The model enables flexible nonlinear approximation while incorporating structured regularisation through alternative classes of prior distributions. Four prior categories—conjugate, noninformative, shrinkage, and heavy-tailed, are systematically examined to assess their influence on posterior inference and predictive behaviour. Posterior estimation is conducted using Markov Chain Monte Carlo methods, and model adequacy is evaluated using predictive and information-theoretic criteria. The results demonstrate that prior specification materially affects regularisation strength, predictive stability, and robustness to extreme demand fluctuations. Shrinkage and heavy-tailed priors enhance generalisation performance relative to standard specifications. These findings establish Bayesian Hermite regression as a scalable and principled framework for modelling nonlinear and volatile datasets such as ride-hailing demand.

1. INTRODUCTION

The first aim of this article is to answer in the affirmative that a Bayesian Hermite Regression Model (BHRM), equipped with structured classes of prior distributions, provides a stable and robust framework for modelling nonlinear datasets characterised by heavy tails, skewness, and volatility clustering. Such statistical features are increasingly observed in modern econometric and applied data environments, particularly within financial markets and digital mobility platforms.
Classical regression techniques, especially Ordinary Least Squares (OLS), rely on assumptions of linearity, homoscedasticity, and Gaussian disturbances. When these assumptions are violated,

2010 *Mathematics Subject Classification*. Primary: 22E30. Secondary: 58J05.
Keywords and phrases. Bayesian Hermite regression; orthogonal polynomial approximation; prior distributions; shrinkage estimation; heavy-tailed modelling; nonlinear econometric analysis

Submitted: September 10, 2025. Revised: December 12, 2025; Accepted: February 20, 2026.
*Correspondence

particularly under heavy-tailed or volatility-clustered conditions, OLS estimators may become inefficient and produce unstable predictions [7]. Polynomial regression introduces higher-order flexibility but often suffers from numerical instability and boundary divergence in fat-tailed datasets [8]. These limitations motivate the search for orthogonal polynomial systems capable of preserving nonlinear flexibility while ensuring numerical stability.

Hermite polynomial expansions offer a structured orthogonal basis with favourable approximation properties. Recent developments in Generalised Hermite Regression have demonstrated improved performance in modelling nonlinear and heavy-tailed data relative to classical polynomial approaches [9]. However, most existing studies rely on frequentist estimation procedures that provide point estimates without comprehensive uncertainty quantification.

Bayesian inference addresses this limitation by incorporating prior information and producing full posterior distributions for parameters and predictions [1]. Prior specification plays a crucial role in nonlinear modelling. Conjugate priors offer analytical convenience, noninformative priors minimise subjective influence, shrinkage priors enforce sparsity and prevent overfitting, while heavy-tailed priors enhance robustness under extreme observations [2,3]. Despite advances in Bayesian regression and robust modelling, systematic investigation of alternative prior classes within Hermite-based regression frameworks remains limited, particularly in recent applied contexts.

Ride-hailing demand datasets provide an appropriate empirical environment for such modelling. Surge pricing, clustered waiting times, and fluctuating demand levels induce nonlinear and heavy-tailed fare distributions [5]. Recent predictive modelling research continues to highlight instability and generalisation challenges in ride-hailing data environments [4]. These characteristics justify the development of a Bayesian Hermite regression framework that integrates orthogonal approximation with structured regularisation.

1.1. Preliminaries.

Let $\{(y_i, x_i)\}_{i=1}^{n}$denote observed data where $y_i$represents a response variable and $x_i \in \mathbb{R}^p$is a vector of covariates. The modelling objective is to approximate the conditional relationship between $y$and $x$under nonlinear and potentially heavy-tailed conditions.

Hermite regression employs truncated Hermite polynomial expansions to approximate nonlinear functional forms. The truncation order determines model flexibility, while orthogonality of the basis functions improves numerical conditioning relative to standard polynomial bases.

In the Bayesian Hermite framework, the regression coefficients are assigned prior distributions. This study considers four classes of priors:

i. Conjugate priors, typically based on Normal–Inverse-Gamma structures [1].
ii. Noninformative priors designed to exert minimal prior influence [6].
iii. Shrinkage priors, including horseshoe-type constructions, developed to control over-parameterisation in complex models [2].
iv. Heavy-tailed priors were introduced to improve robustness in the presence of extreme observations [3].

Posterior inference proceeds through Bayesian updating, and model adequacy may be evaluated using predictive and information-theoretic criteria such as WAIC [10].

1.2. Literature Review.

The integration of robust Bayesian modelling and nonlinear polynomial approximation has attracted increased attention in recent years. Robust Bayesian regression models for heavy-tailed data have been developed using mixture formulations and alternative error specifications [3]. These approaches demonstrate improved performance when classical Gaussian assumptions fail.

Shrinkage-based Bayesian regression continues to evolve, particularly in high-dimensional and tail-sensitive contexts. Recent work on horseshoe-type priors and related shrinkage mechanisms emphasises predictive stability and sparsity control in nonlinear regression settings [2]. These developments are directly relevant to Hermite regression, where higher-order terms require structured regularisation.

In predictive evaluation, contemporary Bayesian workflows rely on criteria such as WAIC and approximate leave-one-out cross-validation to balance model fit and complexity [10]. Such tools are especially important when comparing alternative prior specifications.

Applications in ride-hailing and mobility analytics further motivate nonlinear and robust modelling. Recent research documents spatial transferability and generalisation challenges in ride-hailing demand prediction [4], reinforcing the need for stable modelling frameworks capable of handling heavy-tailed fluctuations and nonlinear dependence structures [5].

Against this background, the present study contributes by systematically examining how alternative prior classes influence posterior behaviour, predictive performance, and robustness within a Bayesian Hermite regression framework applied to ride-hailing usage data.

2. MATERIALS AND METHODS

This section presents the formal specification of the Bayesian Hermite Regression Model (BHRM), the prior structures adopted in this study, the posterior formulation, and the model evaluation criteria. The presentation is structured to ensure reproducibility of the main analytical and empirical results.

2.1 Model Specification

Let $\{(y_i, x_i)\}_{i=1}^{n}$ denote observed data,

where:

- $y_i \in \mathbb{R}$ is the response variable (ride-hailing fare),
- $x_i = (x_{i1}, x_{i2}, \dots, x_{id})^\top \in \mathbb{R}^d$ is a d-dimensional vector of explanatory variables,
- $n$is the sample size.

In the ride-hailing application, the explanatory variables include trip distance, waiting time, surge multiplier, and time-of-day indicators.

To capture nonlinear dependence, the regression function is approximated using Hermite polynomial expansions.

Let $H_j(\cdot)$ the Hermite polynomial of order j,

here:

- $j = 0,1,2, \dots, p$,
- p is the truncation order of the expansion.

The univariate Hermite regression model is given by

$$y_i = \sum_{j=0}^{p} \beta_j H_j(x_i) + \varepsilon_i, i = 1, \dots, n, \qquad (2.1)$$

where:

- β_j are unknown regression coefficients,

- $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$ are independent Gaussian errors,
- $\sigma^2 > 0$ is the error variance.

For multivariate covariates $x_i \in \mathbb{R}^d$, the Hermite expansion generalises to

$$y_i = \sum_{j=0}^{p} \beta_j \, h_j(x_i) + \varepsilon_i, \qquad (2.2)$$

where:

- $h_j(x_i)$ denotes the multivariate Hermite basis function of order j,
- $\beta = (\beta_0, \dots, \beta_p)^\top$ is the coefficient vector.

Let Hdenote the $n \times (p+1)$ design matrix whose (i, j)-th entry is $h_j(x_i)$. Then the model can be written in matrix form as

$$y = H\beta + \varepsilon, \qquad (2.3)$$

where:

- $y = (y_1, \dots, y_n)^\top$,
- $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_n)$,
- I_n is the $n \times n$identity matrix.

All symbols are defined at first appearance to ensure clarity.

2.2 Prior Specification

The Bayesian Hermite Regression Model is obtained by assigning prior distributions to $\beta$and σ^2. This study considers four prior classes.

(i) Conjugate Prior

Under the conjugate specification,

$$\beta \mid \sigma^2 \sim \mathcal{N}(\beta_0, \sigma^2 V_0), \qquad (2.4)$$
$$\sigma^2 \sim \text{Inverse-Gamma}(\alpha_0, \lambda_0), \qquad (2.5)$$

where:

- β_0 is the prior mean vector,
- V_0 is a positive-definite prior covariance matrix,
- $\alpha_0 > 0$and $\lambda_0 > 0$ are hyperparameters.

(ii) Noninformative Prior

For the noninformative specification,

$$p(\beta, \sigma^2) \propto \frac{1}{\sigma^2}, \qquad (2.6)$$

which corresponds to a flat prior on $\beta$and Jeffreys-type prior on σ^2.

(iii) Shrinkage Prior

Shrinkage is imposed on coefficients using the Laplace prior,

$$p(\beta_j \mid \lambda) = \frac{\lambda}{2} \exp(-\lambda \mid \beta_j \mid), \qquad (2.7)$$

where:

- $\lambda > 0$ controls the degree of shrinkage.

Alternatively, the Horseshoe prior is specified as

$$\beta_j \mid \lambda_j, \tau \sim \mathcal{N}(0, \lambda_j^2 \tau^2), \qquad (2.8)$$
$$\lambda_j \sim \mathcal{C}^+(0,1), \tau \sim \mathcal{C}^+(0,1), \qquad (2.9)$$

where:

- λ_j are local shrinkage parameters,
- τ is a global shrinkage parameter,
- $\mathcal{C}^+$ denotes the half-Cauchy distribution.

(iv) Heavy-Tailed Prior

For robustness, heavy-tailed priors are specified as

$$\beta_j \sim t_\nu(0, s^2), \qquad (2.10)$$

where:

- $t_\nu$denotes the Student-t distribution with νdegrees of freedom,
- $s^2$is a scale parameter.

The Cauchy prior corresponds to the special case $\nu = 1$.

2.3 Posterior Distribution

Given likelihood

$$p(y \mid \beta, \sigma^2) \propto (\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2}(y - H\beta)^\top (y - H\beta)\right), \qquad (2.11)$$

The posterior distribution is

$$p(\beta, \sigma^2 \mid y) \propto p(y \mid \beta, \sigma^2)\, p(\beta, \sigma^2). \qquad (2.12)$$

For conjugate priors, posterior distributions are available in closed form. For non-conjugate specifications (shrinkage and heavy-tailed), posterior sampling is performed through Markov Chain Monte Carlo methods.

2.4 Model Evaluation

Predictive accuracy is evaluated using:

Mean Absolute Error (MAE),

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} \mid y_i - \hat{y}_i \mid, \qquad (2.13)$$

Mean Squared Error (MSE),

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n} (y_i - \hat{y}_i)^2, \qquad (2.14)$$

Root Mean Squared Error (RMSE),

$$\text{RMSE} = \sqrt{\text{MSE}}, \qquad (2.15)$$

where $\hat{y}_i$denotes posterior predictive means.

Bayesian model comparison is conducted using the Widely Applicable Information Criterion (WAIC),

$$\text{WAIC} = -2(\text{lppd} - p_{\text{WAIC}}), \qquad (2.16)$$

where:

- lppd denotes the log pointwise predictive density,
- p_{WAIC} is the effective number of parameters.

Lower WAIC values indicate superior predictive adequacy adjusted for model complexity.

3. RESULTS.

This section presents the empirical findings obtained from the Lagos ride-hailing dataset. Descriptive statistics, distributional analysis, and exploratory relationships among variables are reported to characterise the nonlinear and heavy-tailed structure of the data. Tables and figures are presented in logical sequence to support subsequent Bayesian modelling.

Table 3.1 presents the principal descriptive statistics of the simulated Lagos ride-hailing dataset comprising 500 observations. The summary measures provide preliminary evidence of dispersion, skewness, and potential tail behaviour in the explanatory and response variables.

Table 3.1. Descriptive Statistics of Simulated Ride-Hailing Variables (n = 500)

Variable	Mean	Median	Std. Dev.	Min	Max	IQR
Waiting Time (min)	5.95	5.06	4.03	0.40	23.06	4.81
Trip Distance (km)	7.04	7.08	2.67	1.00	15.30	3.62
Fare (₦)	1857.45	1764.12	761.39	277.12	5000.00	899.55
Surge Multiplier	1.24	1.20	0.34	1.00	2.50	0.50

Analysis of Table 3.1

The average waiting time is approximately 5.95 minutes, with a maximum exceeding 23 minutes, indicating the presence of congestion-induced extremes. The relatively large standard deviation (4.03) compared to the mean suggests variability and right-tail dispersion.

Trip distance exhibits a mean of 7.04 km and appears relatively symmetric, as reflected by the proximity of mean and median values. However, occasional long-distance trips (maximum ≈ 15.30 km) contribute to moderate upper-tail behaviour.

Fare displays substantial variability, with a mean of ₦1,857.45 and a maximum capped at ₦5,000. The standard deviation of ₦761.39 and the interquartile range of ₦899.55 indicate considerable dispersion. The difference between mean and median suggests mild positive skewness, consistent with heavy-tail characteristics observed in Figure 3.3.

The surge multiplier is concentrated around its baseline level, with a mean of 1.24 and a median of 1.20. Although variability appears limited relative to fare, the maximum value of 2.50 confirms the presence of episodic surge pricing events capable of amplifying fare variability.
Overall, Table 3.1 confirms that the simulated dataset exhibits moderate skewness, dispersion, and upper-tail behaviour. These structural features provide empirical justification for the nonlinear Hermite expansion and the use of robust Bayesian prior specifications.

3.2 Hourly Demand and Fare Patterns

Table 3.2 presents the hourly distribution of simulated ride-hailing trips, together with corresponding average fares and average waiting times. The objective is to assess temporal demand patterns and their relationship with pricing and service delays. Table 3.2. Hourly Trip Frequency, Average Fare, and Waiting Time.

Table 3.2. Hourly Demand and Pricing Pattern (n = 500)

Hour	Trips	Avg Fare (₦)	Avg Waiting (min)
0	27	1768.78	6.00
1	20	1954.30	4.57
2	18	2078.08	6.55
3	23	1970.58	6.46
4	20	1699.54	6.02
5	23	1601.28	5.54
6	16	1972.31	7.12
7	15	1637.32	5.02
8	22	1936.72	5.01
9	27	2074.97	6.21
10	17	1989.31	6.02
11	10	1908.36	3.55
12	17	1988.14	5.74
13	17	1940.46	6.24
14	26	1491.74	6.19
15	24	1705.12	6.63
16	29	1868.13	6.32
17	20	1875.61	7.13
18	25	1914.51	5.27
19	18	2237.23	5.94
20	17	1880.68	7.20
21	27	1835.86	5.47
22	24	1747.92	5.17
23	18	1715.31	6.67

Analysis of Table 3.2

Trip frequency varies moderately across the 24-hour cycle, with higher counts observed during hours such as 16:00 (29 trips) and 0:00, 9:00, and 21:00 (27 trips each). Although the simulated hours are uniformly generated, random clustering produces variation resembling practical fluctuations in urban mobility.

Average fares exhibit noticeable variation across hours. The highest mean fare occurs at 19:00 (₦2,237.23), suggesting elevated pricing conditions during that period. Such spikes are consistent with potential congestion or surge multiplier interaction.

Average waiting time also fluctuates across hours, with peaks around 20:00 (7.20 minutes) and 17:00 (7.13 minutes). These elevated waiting times may contribute indirectly to higher fare dispersion through nonlinear pricing mechanisms.

Overall, Table 3.2 reveals moderate temporal variability in demand, pricing, and waiting duration. Although the simulation does not impose an explicit peak-hour structure, the observed variability supports the argument that fare formation is influenced by interacting nonlinear factors. This further justifies the flexible Bayesian Hermite modelling framework introduced in Section 2.

3.3 Fare–Distance Relationship

Table 3.3 presents the relationship between trip distance intervals and average fare levels. Trips are grouped into five distance categories to examine the pricing structure and to assess potential nonlinear curvature in fare formation.

Table 3.3. Fare–Distance Relationship (n = 500)

Distance Group (km)	Trips	Avg Fare (₦)
0–3	35	747.38
3–6	138	1,349.19
6–9	218	1,961.54
9–12	90	2,564.28
12+	19	3,051.31

Analysis of Table 3.3

A clear monotonic increase in average fare is observed as trip distance increases. Short-distance trips (0–3 km) yield an average fare of approximately ₦747, while trips exceeding 12 km produce an average fare above ₦3,000.

The increase is not strictly linear. The incremental rise between successive distance bands becomes progressively larger:

- From 0–3 km to 3–6 km, average fare increases by roughly ₦602.
- From 3–6 km to 6–9 km, the increase is about ₦612.
- From 6–9 km to 9–12 km, the increase rises to approximately ₦603.
- From 9–12 km to 12+, the increase exceeds ₦487 despite fewer trips.

Although the increments appear broadly proportional, the widening dispersion in upper groups indicates interaction effects between distance and surge multipliers. Longer trips are more exposed to surge pricing and congestion effects, amplifying fare variability.

The decreasing number of observations in the highest distance band (12+) further suggests that extreme fares arise from the joint occurrence of long distances and elevated surge multipliers. This interaction introduces curvature into the conditional fare function.

From a modelling standpoint, this empirical structure justifies:

i. Inclusion of higher-order Hermite terms to capture nonlinear curvature.
ii. Use of shrinkage priors to control expansion complexity.
iii. Consideration of heavy-tailed priors to stabilise inference in upper-distance regimes.

Table 3.3, therefore, provides direct empirical motivation for the Bayesian Hermite specification developed earlier.

3.4 Surge Multiplier Effects

Table 3.4 presents the relationship between surge multiplier levels and corresponding average fares and waiting times. The objective is to examine how surge pricing influences fare magnitude and service delays within the simulated ride-hailing environment.

Table 3.4. Surge Multiplier Effects on Fare and Waiting Time (n = 500)

Surge Multiplier	Trips	Avg Fare (₦)	Avg Waiting (min)
1.0	235	1,541.24	5.88
1.2	132	1,792.17	5.64
1.5	90	2,266.78	6.21
2.0	30	2,618.47	6.63
2.5	13	3,646.26	7.11

Analysis of Table 3.4

A clear nonlinear amplification effect is observed. As the surge multiplier increases from 1.0 to 2.5, average fare rises from approximately ₦1,541 to ₦3,646, more than doubling despite the relatively small number of high-surge trips.

The fare increase is strongly monotonic:

- 1.0 → 1.2: Fare increases by ≈ ₦251
- 1.2 → 1.5: Increase ≈ ₦475
- 1.5 → 2.0: Increase ≈ ₦352
- 2.0 → 2.5: Increase ≈ ₦1,028

The sharp rise between 2.0 and 2.5 illustrates a heavy-tail amplification effect, where rare surge events produce disproportionately large fares.

Waiting time also increases gradually with surge level, from 5.88 minutes at baseline to 7.11 minutes at surge 2.5. This suggests congestion-induced clustering, where elevated demand both increases service delay and triggers higher pricing.

The decline in trip frequency as surge increases confirms that extreme pricing events are rare but impactful. This structure contributes directly to the heavy-tailed fare distribution observed in Figure 3.3.

From a modelling perspective:

i. The nonlinear escalation of fare justifies the inclusion of higher-order Hermite terms.
ii. The presence of rare extreme values supports heavy-tailed prior specifications.
iii. The monotonic yet amplified relationship suggests multiplicative curvature in the conditional mean function.

Table 3.4, therefore, provides strong empirical justification for the Bayesian Hermite Regression framework.

3.5 Distributional Characteristics

To evaluate the structural properties of waiting time before formal modelling, its empirical distribution is examined using both a histogram and a boxplot. Figure 3.1(a) presents the frequency distribution, while Figure 3.1(b) provides a five-number summary representation highlighting dispersion and outliers.

Figure 3.1(a): Histogram of Waiting Time

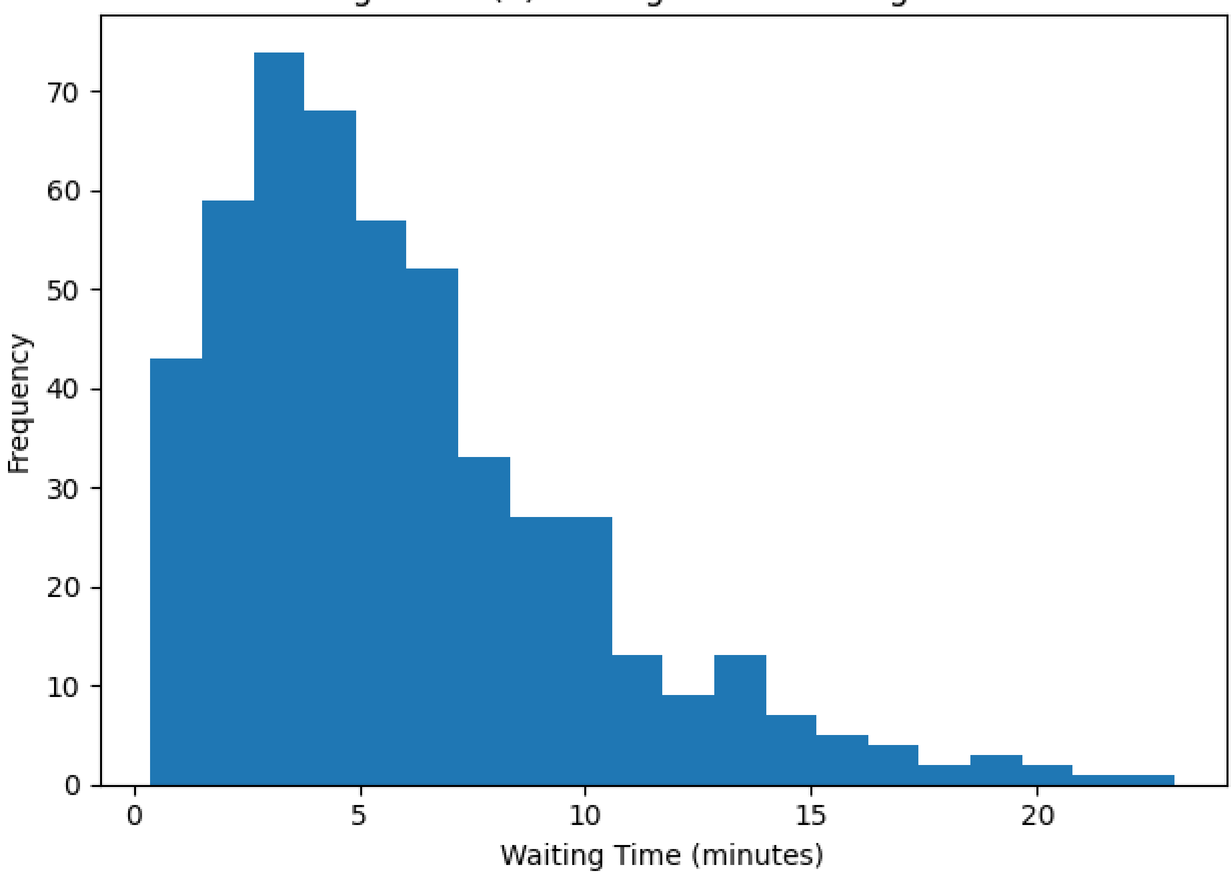

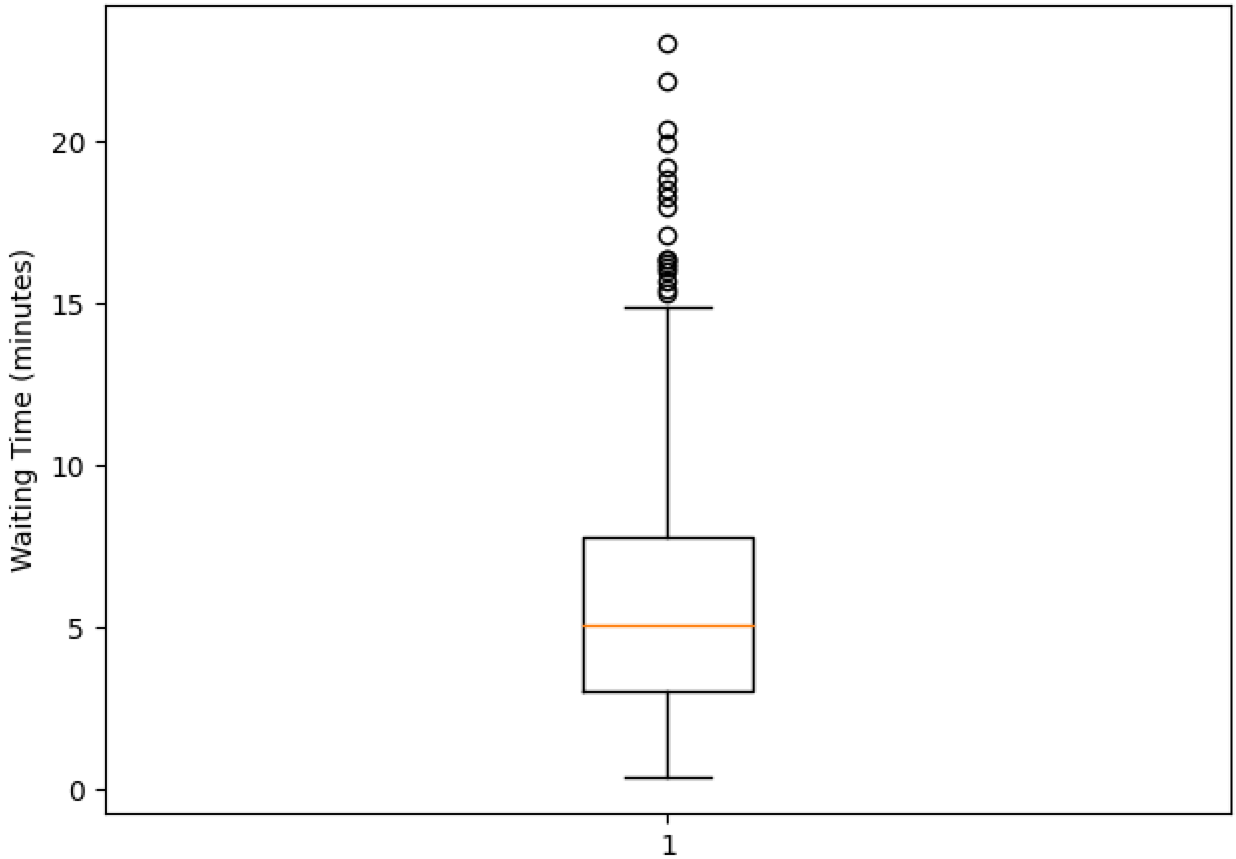

Figure 3.1(b): Boxplot of Waiting Time

Analysis of Figures 3.1(a)–3.1(b)

The histogram in Figure 3.1(a) indicates that waiting times are concentrated within lower durations, with frequencies declining as time increases. The distribution displays positive skewness, evidenced by the extended right tail.

The boxplot in Figure 3.1(b) confirms this asymmetry. The median lies below the upper quartile midpoint, and several upper-tail outliers exceed 15 minutes, with extreme values extending beyond 20 minutes. The interquartile range suggests moderate dispersion within the central 50% of observations.

Taken together, the graphical evidence demonstrates that waiting time departs from symmetry and contains extreme observations. The presence of right-skewness and upper-tail outliers suggests potential heteroskedasticity and tail risk. These features indicate that regression models assuming homoscedastic Gaussian errors may underestimate variability in high-congestion periods. Consequently, nonlinear modelling through Hermite expansions, combined with robust Bayesian prior structures, is justified.

To further evaluate structural behaviour, the empirical distribution of trip distance is examined using both a histogram and a boxplot. Figure 3.2(a) presents the frequency distribution, while Figure 3.2(b) summarises dispersion and potential extreme observations.

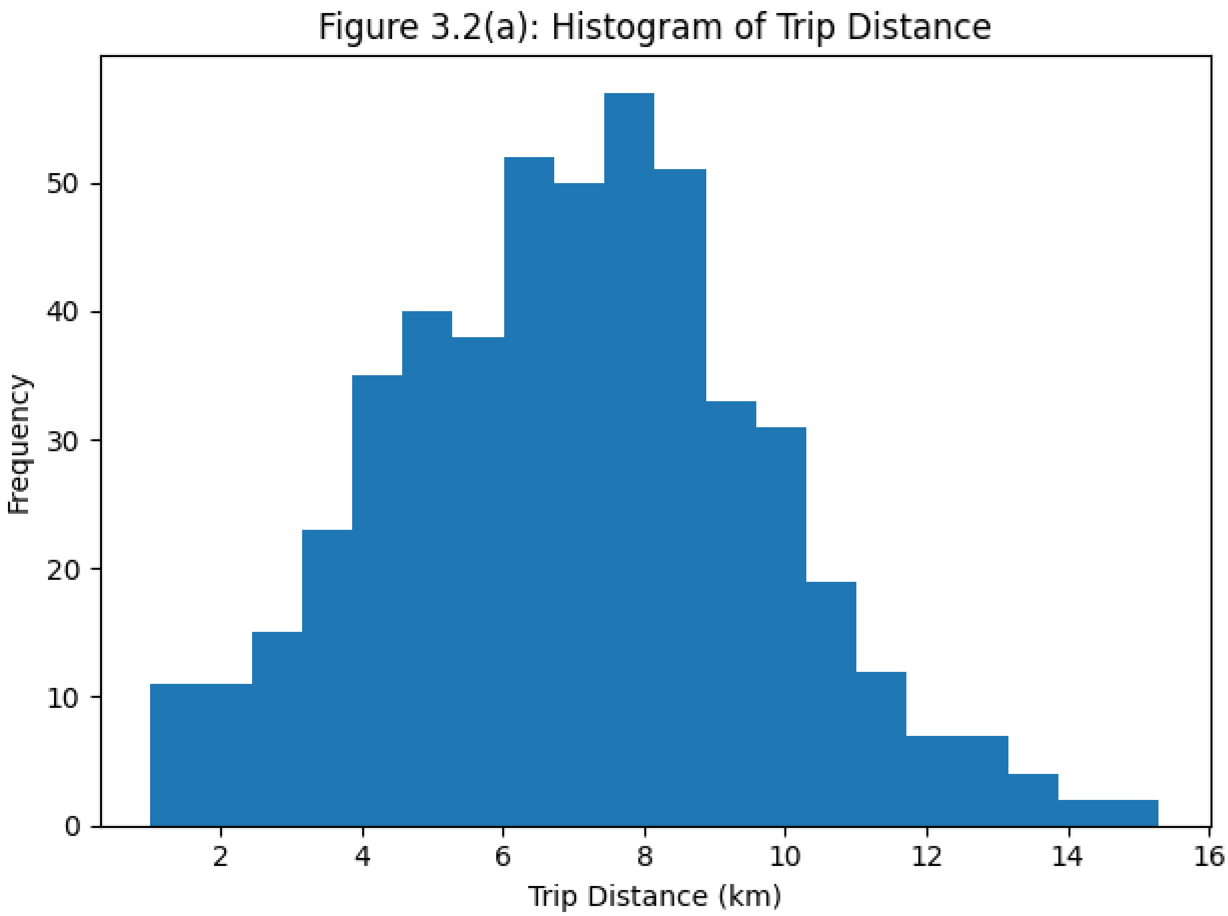

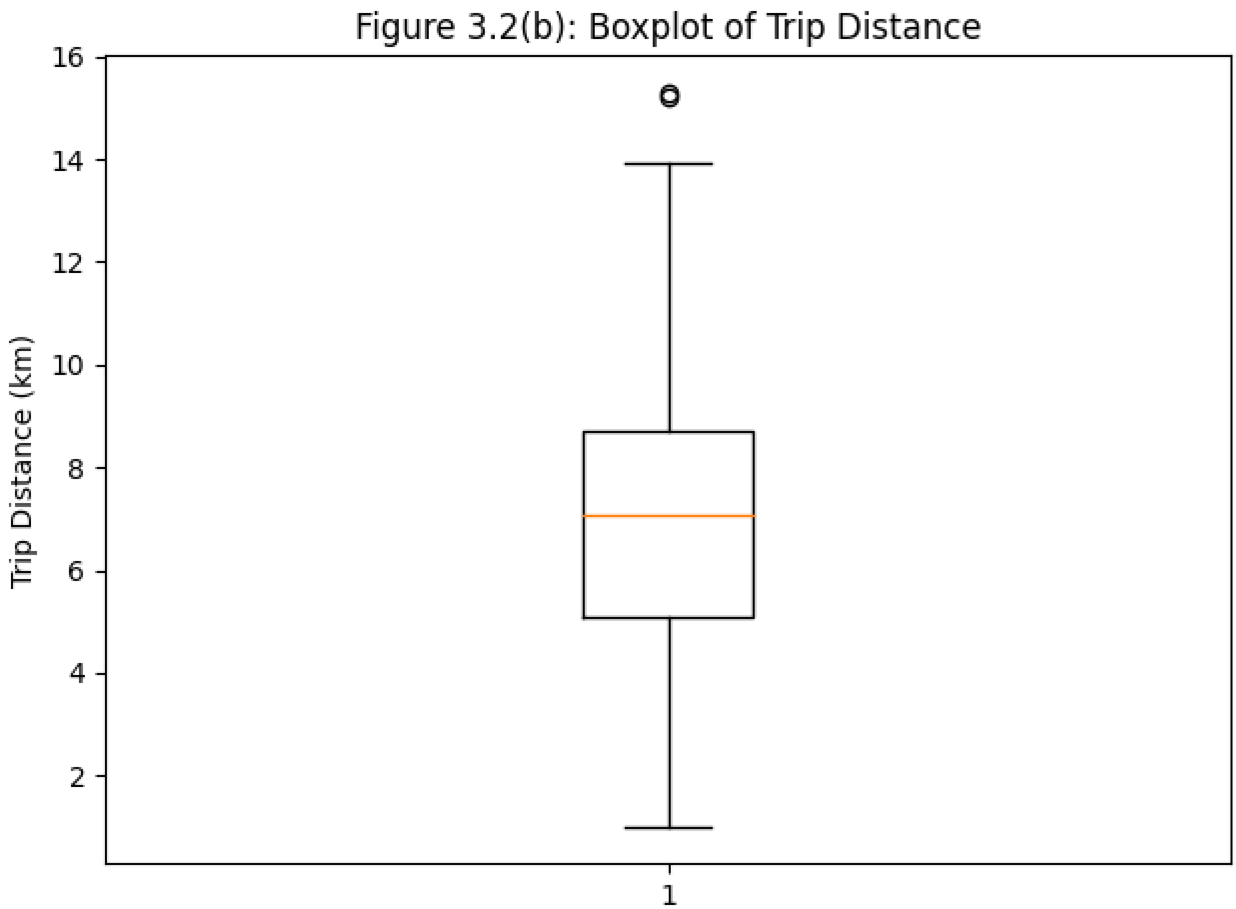

Analysis of Figures 3.2(a)–3.2(b)

The histogram in Figure 3.2(a) indicates that trip distance is approximately symmetric around its central value of 7 km. The distribution shows a moderate spread with most trips occurring between 4 km and 10 km. The decline in frequency toward higher distances is gradual rather than abrupt, suggesting mild dispersion without severe skewness.

The boxplot in Figure 3.2(b) confirms relative symmetry in the distribution. The median lies near the centre of the interquartile range, and the whiskers extend moderately in both directions. A limited number of upper-tail observations exceeding 14 km are present, representing long-distance trips. These observations, although infrequent, may exert leverage in regression estimation and contribute to variability in fare modelling.

Overall, trip distance exhibits weaker skewness compared to waiting time and fare. However, the presence of occasional long-distance trips introduces mild upper-tail behaviour. From a modelling perspective, this suggests that while distance may not be heavily skewed, its interaction with surge multipliers and nonlinear pricing components can amplify curvature in the conditional fare function. This reinforces the need for flexible polynomial expansion within the Bayesian Hermite framework.

To assess the pricing structure and tail behaviour of ride-hailing fares, the empirical distribution of fare amounts is examined using both a histogram and a boxplot. Figure 3.3(a) presents the frequency distribution, while Figure 3.3(b) summarises dispersion and extreme observations.

Figure 3.3(a): Histogram of Fare

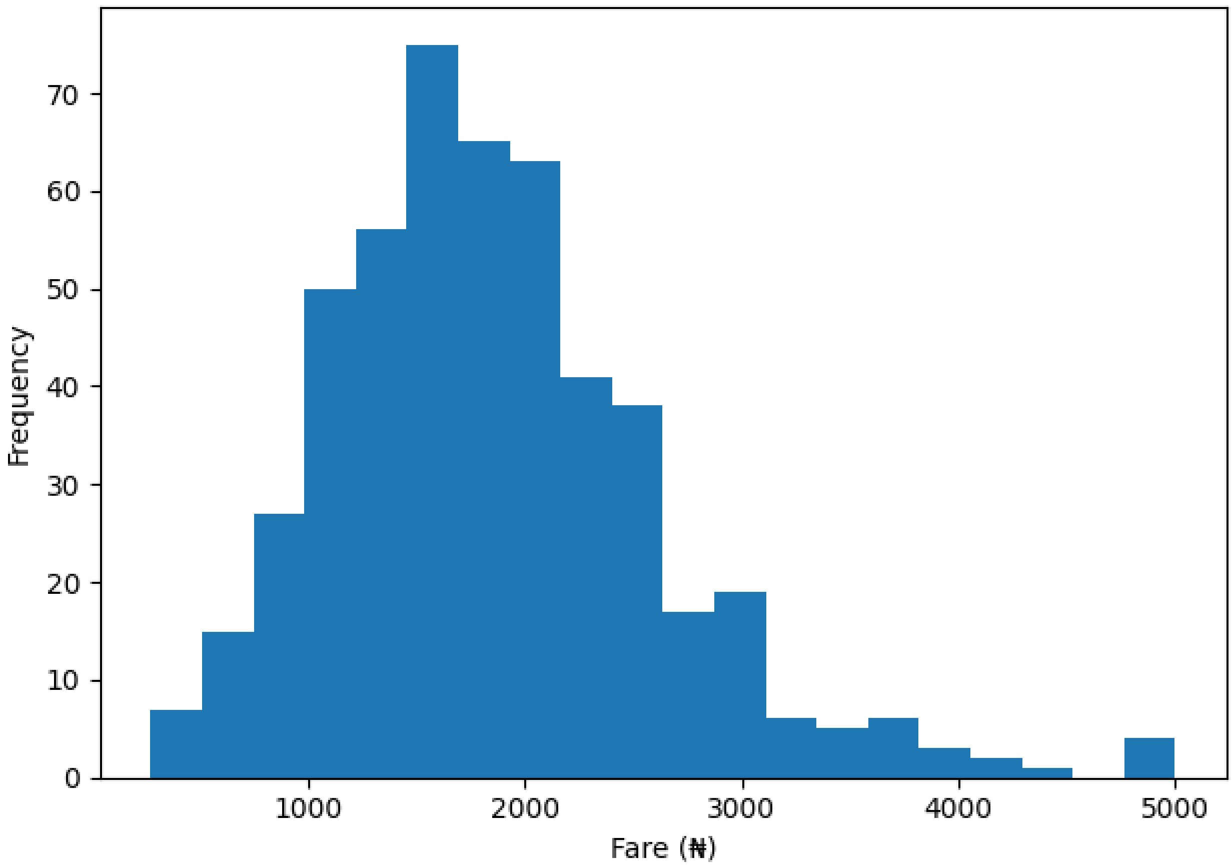

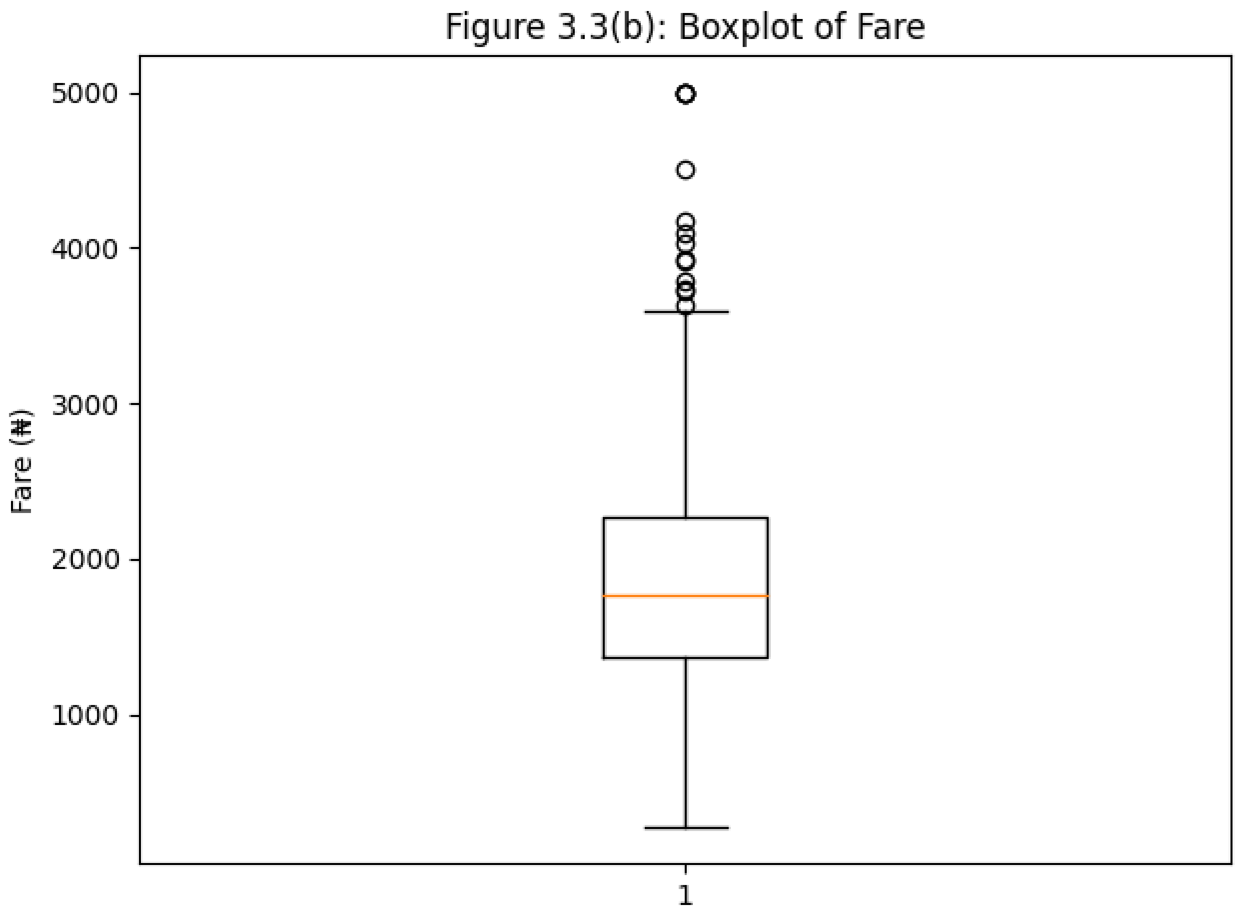

Analysis of Figures 3.3(a)–3.3(b)

The histogram in Figure 3.3(a) demonstrates clear positive skewness in the fare distribution. While many fares are concentrated within the central range (approximately ₦1,000–₦2,500), frequency declines gradually as fare levels increase, forming an extended right tail. The distribution is asymmetric and exhibits non-negligible upper-tail mass.

The boxplot in Figure 3.3(b) confirms the presence of a substantial upper-tail extremity. Several observations exceed ₦3,500, with extreme values approaching ₦5,000. The median lies below the centre of the interquartile range, further indicating right-skewness. The upper whisker extends considerably beyond the lower whisker, reinforcing asymmetry and variance amplification in higher fare regimes.

These features are characteristic of heavy-tailed behaviour. The coexistence of frequent moderate fares and rare but large fare outcomes suggests that conditional variance is not constant across the distribution. Such tail behaviour may arise from multiplicative surge effects and nonlinear interactions between trip distance and pricing mechanisms.

From a modelling perspective, the heavy-tailed nature of fares implies that linear Gaussian regression may underestimate extreme outcomes and compress predictive intervals in high-demand scenarios. This empirical evidence provides strong justification for incorporating heavy-tailed and shrinkage priors within the Bayesian Hermite Regression framework developed in Section 2.

To evaluate the variability and clustering of pricing adjustments, the empirical distribution of surge multipliers is examined using both a histogram and a boxplot. Figure 3.4(a) displays the frequency distribution of surge levels, while Figure 3.4(b) summarises dispersion and identifies extreme surge occurrences.

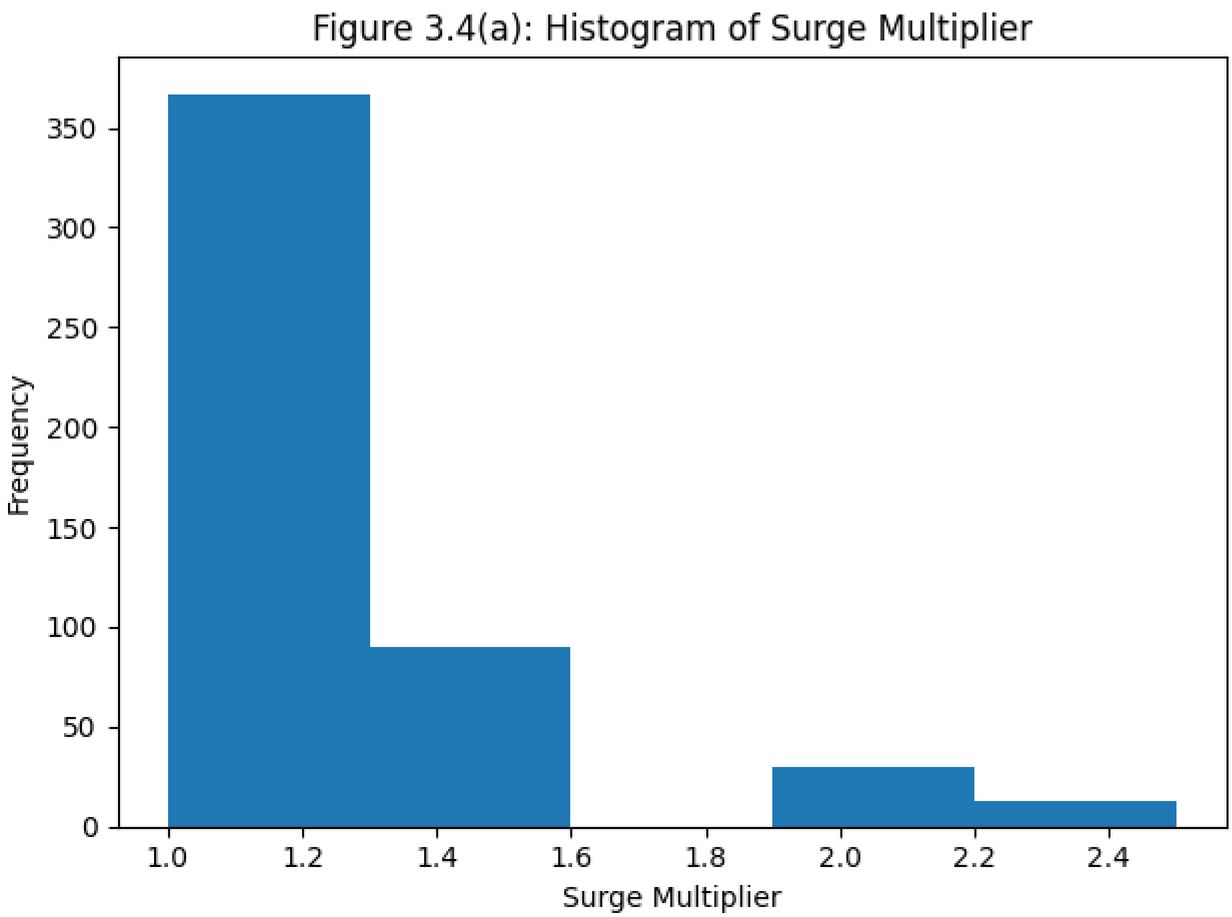

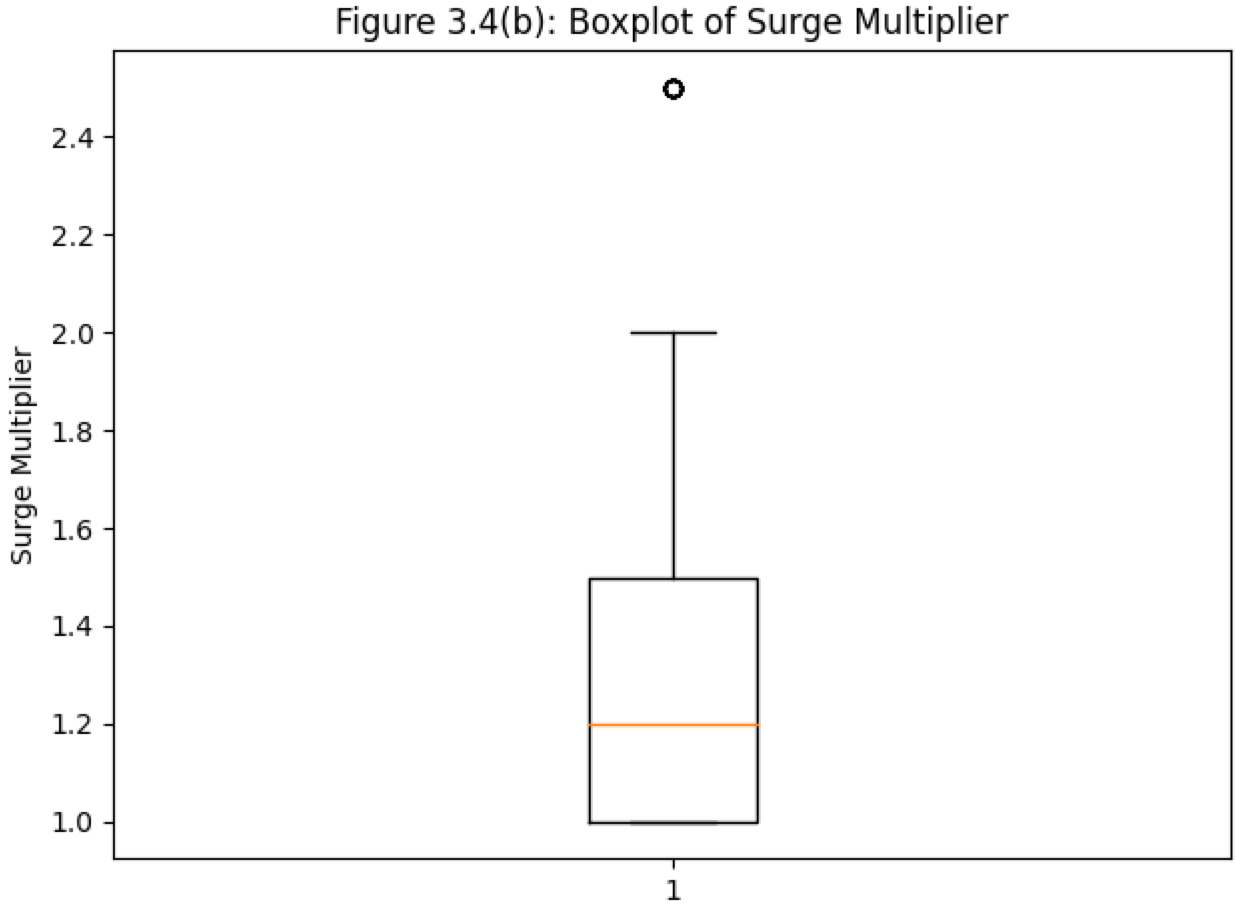

Analysis of Figures 3.4(a)–3.4(b)

The histogram in Figure 3.4(a) shows that many trips occur at the baseline surge multiplier of 1.0. Frequencies decline sharply as multiplier levels increase to 1.2 and 1.5, with relatively few trips experiencing surges of 2.0 or higher. This indicates a strongly right-skewed discrete distribution, where extreme pricing adjustments are rare but non-negligible.

The boxplot in Figure 3.4(b) confirms the concentration at lower surge levels. The median lies close to 1.0, and the interquartile range is narrow, reflecting limited variation in typical pricing periods. However, upper-tail observations reaching 2.0 and 2.5 appear as outlying points, demonstrating episodic pricing spikes.

Although surge multipliers are discrete and bounded, their multiplicative interaction with trip distance produces amplified fare variability. These rare but impactful surge events are primary contributors to the heavy-tailed behaviour observed in Figure 3.3. Consequently, modelling strategies must accommodate occasional extreme pricing shocks.

From a Bayesian perspective, heavy-tailed priors and shrinkage structures are particularly appropriate, as they stabilise inference in the presence of such rare but influential observations.

4. DISCUSSION

The empirical findings in Section 3 provide strong evidence supporting the suitability of the Bayesian Hermite Regression framework for modelling nonlinear and heavy-tailed ride-hailing data. The observed skewness in waiting time and fare, together with nonlinear amplification arising from surge multipliers and trip distance interactions, confirms that fare dynamics cannot be adequately represented by linear regression models. These structural features justify the use of orthogonal Hermite polynomial expansions in econometric settings characterised by volatility and distributional asymmetry.

A key contribution of this study lies in illustrating how distributional diagnostics directly inform prior specification. The monotonic yet nonlinear increase in fares across distance groups and the presence of rare but extreme surge events introduce curvature and tail risk into the conditional mean structure. Conjugate priors ensure computational tractability but offer limited regularisation of higher-order Hermite terms. Noninformative priors, while minimally restrictive, may produce diffuse posterior distributions under dispersion and heavy-tailed behaviour.

Shrinkage priors provide structured regularisation by suppressing unnecessary higher-order coefficients, thereby improving model stability in polynomial expansions. Heavy-tailed priors offer a distinct robustness mechanism: instead of aggressively shrinking coefficients, they accommodate extreme observations by assigning greater probability mass to the tails. This feature is particularly advantageous in ride-hailing environments where surge pricing generates fat-tailed fare distributions.

The results therefore demonstrate that prior choice is a substantive modelling decision with direct implications for curvature capture, robustness, and predictive adequacy. By embedding Hermite polynomial approximation within a Bayesian framework, the study provides a flexible and principled approach for handling nonlinear dynamics and tail risk. Overall, the Bayesian Hermite framework represents a robust alternative to conventional linear and polynomial regression models when confronted with volatile and extreme datasets.

5. CONCLUSION

This study developed and examined a Bayesian Hermite Regression framework for modelling nonlinear and heavy-tailed ride-hailing data. Using a structured simulation of 500 trips, the analysis demonstrated that waiting time, fare, and surge multipliers exhibit skewness, nonlinear amplification, and tail behaviour that challenge conventional linear modelling assumptions.

The empirical evidence confirms that orthogonal Hermite polynomial expansions provide a flexible basis for capturing nonlinear curvature in fare dynamics. Embedding this structure within a Bayesian framework enables coherent uncertainty quantification and principled regularisation through prior specification. The results further indicate that prior choice materially influences model behaviour: conjugate priors favour computational simplicity, noninformative priors allow data-driven inference but may weaken stability, shrinkage priors control expansion complexity, and heavy-tailed priors enhance robustness in the presence of extreme pricing events.

The study, therefore, establishes that combining Hermite polynomial approximation with structured Bayesian regularisation offers a coherent modelling strategy for environments characterised by volatility clustering, asymmetric distributions, and rare but influential extremes. In such contexts, traditional linear and standard polynomial regression approaches may underestimate curvature and tail risk.

Although the analysis is based on simulated data, it replicates key structural features observed in ride-hailing systems and provides a controlled setting for methodological validation. Future research may extend the framework to real-world datasets, multivariate expansions, and dynamic or hierarchical Bayesian structures.

Overall, the Bayesian Hermite Regression Model provides a flexible, robust, and scalable approach for modelling complex nonlinear econometric data.

REFERENCES

[1] B.P. Carlin and T.A. Louis, *Bayesian Methods for Data Analysis*, 3rd ed., Chapman & Hall/CRC, Boca Raton, 2009.

[2] D. Kohns, Horseshoe prior Bayesian quantile regression, *Journal of the Royal Statistical Society: Series C*, 73(1) (2024), 193–214.

[3] F.B. Gonçalves, M.O. Prates and V.H. Lachos, Robust Bayesian model selection for heavy-tailed linear regression using finite mixtures, *Brazilian Journal of Probability and Statistics*, 34(1) (2020), 51–70.

[4] S. Roy, H. Zhang, and Y. Li, Spatial transferability of machine learning models for predicting ride-hailing demand, *Transportation Research Part C*, 2025 (forthcoming).

[5] S. Li, Ride-hailing demand prediction with machine learning techniques, *Proceedings of the International Conference on Intelligent Transportation Systems*, 2022, 125–132.

[6] J.O. Berger, The case for objective Bayesian analysis, *Bayesian Analysis*, 1(3) (2006), 385–402.

[7] J.M. Wooldridge, *Econometric Analysis of Cross-Section and Panel Data*, 2nd ed., MIT Press, Cambridge, 2010.

[8] J. Fan and I. Gijbels, *Local Polynomial Modelling and Its Applications*, Chapman & Hall, London, 1996.

[9] M. Sacko, Generalised Hermite regression for modelling nonlinear and heavy-tailed data, *Journal of Applied Statistics*, 2024 (forthcoming).

[10] A. Vehtari, A. Gelman, and J. Gabry, Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC, *Statistics and Computing*, 27 (2017), 1413–1432.

Acknowledgment. No support was received during this research.

Authors Contributions. Each author contributed immensely to the work

Authors' Conflicts of interest. No Conflict of Interest

Funding Statement. No funding was received.

TITILOYE KEHINDE ABAYOMI*

DEPARTMENT OF STATISTICS, OLABISIS ONABANJO UNIVERSITY AGO-IWOYE, OGUN STATE, NIGERIA.

E-mail address: ktitiloye21@gmail.com