

Bayesian Hierarchical Modeling for Inflation Forecasting in Emerging Economies

Ameer Musa Imran Alhseeni *

Planning Department, General Directorate of Qadisiyah Education, Ministry of Education

*Corresponding author E-mail: amer.alhussani@gmail.com

Abstract

This paper develops a Bayesian hierarchical model to forecast inflation rates in emerging economies, incorporating structural heterogeneity across countries. Unlike traditional models, the proposed approach allows for both country-specific dynamics and global information sharing. The model is implemented using Hamiltonian Monte Carlo methods and evaluated through a simulation study and real data analysis involving Iraq, Egypt, Turkey, and India. Key macroeconomic predictors include exchange rates, interest rates, and oil prices. The results demonstrate that the hierarchical model outperforms conventional approaches such as ARIMA and VAR in terms of forecast accuracy, parameter stability, and uncertainty quantification. This highlights the model's potential for more informed macroeconomic planning in volatile and data-constrained environments.

Keywords: Bayesian Hierarchical Model; Inflation Forecasting; Emerging Economies; Hamiltonian Monte Carlo; Macroeconomic Uncertainty.

1. Introduction

Forecasting inflation remains a critical policy objective, especially in emerging economies facing structural instability, commodity reliance, and volatile financial systems (Cheng & Phillips, 2019). Traditional models such as ARIMA or VAR often rely on assumptions of cross-country homogeneity, failing to capture institutional, monetary, and geopolitical variations that drive inflationary behavior (Canova, 2007). These limitations have encouraged the adoption of more flexible frameworks that accommodate cross-national heterogeneity while still leveraging shared information.

Bayesian hierarchical models provide a coherent statistical solution by incorporating country-specific parameters within a global structure, allowing partial pooling of information (Gelman et al., 2013). This enables more reliable inference in data-scarce settings and improves forecast stability. Moreover, the Bayesian approach facilitates uncertainty quantification and model calibration through posterior distributions (Ghosh et al., 2020).

This study introduces a two-level Bayesian hierarchical regression model to forecast inflation across selected emerging economies, specifically Iraq, Egypt, Turkey, and India, using macroeconomic variables such as exchange rates, interest rates, and oil prices. The model is implemented using Hamiltonian Monte Carlo (HMC) via the Stan framework, which provides efficient posterior sampling (Carpenter et al., 2017). The model's predictive accuracy is evaluated through simulation and real data analysis.

By benchmarking against classical models such as ARIMA and VAR, this research emphasizes the improved accuracy, interpretability, and generalizability of hierarchical Bayesian models in multi-country macroeconomic forecasting.

2. Theoretical background

2.1. Inflation dynamics in emerging markets

Inflation in emerging markets is influenced by a complex interplay of global shocks and local structural conditions. One of the most prominent external drivers is oil price volatility, which directly affects both importing and exporting countries. For instance, cointegration analysis in Ghana confirmed a long-run positive relationship between oil prices and inflation, suggesting that increases in oil prices lead to sustained inflationary pressure (Amoako & Adusei, 2021). In oil-dependent economies like Russia, fluctuations in global oil prices have been shown to significantly affect domestic inflation through their impact on government revenues, exchange rates, and monetary conditions (Bagirov, 2022).

Another critical factor is the exchange rate pass-through (ERPT) effect, whereby changes in the nominal exchange rate influence domestic prices. Evidence from a European Central Bank study across 28 emerging economies showed that ERPT tends to be larger and more persistent in countries with weaker monetary frameworks and higher baseline inflation (ECB, 2008). This aligns with Taylor's hypothesis that countries with high inflation levels are more susceptible to exchange rate-induced price shifts. However, the magnitude of this pass-through varies depending on the exchange rate regime, trade structure, and the degree of central bank independence. Recent research in

Bangladesh found that while ERPT to consumer prices exists, it is less pronounced in the short run, with money supply and global commodity prices playing more dominant roles (Rahman & Jahan, 2023).

Inflation expectations also play a vital role in shaping inflation dynamics. In economies where the central bank lacks strong credibility, inflation expectations tend to be unanchored, diminishing the effectiveness of conventional monetary policy tools. A study on India found that changes in the real exchange rate were closely associated with fluctuations in household inflation expectations, emphasizing the limitations of interest rate adjustments in such settings (Banerjee et al., 2024).

Structural issues further complicate inflation control in emerging markets. Low productivity, ineffective subsidy mechanisms, and high levels of informality in the labor market reduce supply responsiveness, making inflation more persistent in response to demand or cost shocks. According to IMF assessments, countries with weak institutional capacity exhibit slower adjustment mechanisms to inflationary pressures, especially during global crises (IMF, 2023). These observations highlight the inadequacy of using uniform models across heterogeneous economies and underline the need for modeling frameworks that can capture both shared influences and country-specific responses.

2.2. Bayesian hierarchical models

Bayesian hierarchical models (BHMs) offer a flexible and powerful statistical framework for modeling structured data, particularly when observations are nested within higher-level groups. In the context of inflation forecasting for emerging economies, BHMs are ideal for modeling both shared global dynamics and country-specific effects. These models overcome the limitations of conventional approaches by incorporating partial pooling, allowing individual countries to have their inflation determinants while borrowing strength from the entire group.

Let y_{it} denote the inflation rate in the country i at time t , and let $x_{it} = (x_{1it}, x_{2it}, \dots, x_{kit})$ represent k predictors such as oil price, interest rate, and exchange rate. The level-1 (country-level) regression is specified as:

$$y_{it} \sim N(\alpha_i + \beta_i^T x_{it}, \sigma^2) \quad (1)$$

Where α_i is the country-specific intercept, $\beta_i \in R^k$ is a vector of country-specific regression coefficients, and σ^2 is the observation-level variance assumed common across countries.

At the second level, the parameters (α_i, β_i) are assumed to be drawn from a group-level distribution:

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_\alpha \\ \mu_\beta \end{pmatrix}, \Sigma\right) \quad (2)$$

Here, μ_α and μ_β represent the global means, and Σ is the covariance matrix that captures the variation across countries. To complete the Bayesian specification, prior distributions are placed on the hyperparameters. A typical choice is:

$$\mu_\alpha \sim N(0, 10^2), \mu_\beta \sim N(0, 10^2 I_k), \Sigma \sim \text{LKJ} \quad (2)$$

This setup allows for weakly informative priors, encouraging shrinkage but not enforcing overly strong assumptions. The LKJ prior on Σ promotes reasonable correlation structures among parameters across countries. One major strength of BHMs lies in their capacity to improve estimation in small-sample settings common in developing economies by sharing information across groups. This "borrowing strength" approach improves parameter stability and predictive performance, especially when time series are short or missing. Gelman and Hill (2007) and Gelman et al. (2013) emphasize that hierarchical modeling yields more robust estimates than fully pooled or separate models.

Recent applications have confirmed the practical advantages of this framework. Carriero et al. (2019) used hierarchical Bayesian VARs to improve regional inflation forecasts in the U.S., while Loddo et al. (2022) modeled inflation persistence across Eurozone countries, capturing country-level variation in response to monetary shocks.

2.3. Estimation via Hamiltonian Monte Carlo

Bayesian hierarchical models often involve complex posterior distributions that are analytically intractable due to their high dimensionality and strong parameter correlations. Traditional MCMC methods like Metropolis-Hastings or Gibbs sampling can suffer from poor mixing and inefficiency in such contexts. Hamiltonian Monte Carlo (HMC), a gradient-based sampling technique, provides a more efficient alternative by using the geometry of the posterior distribution to guide the sampling process (Neal, 2011).

HMC introduces an auxiliary momentum variable and constructs a Hamiltonian function that represents the total energy of a hypothetical particle moving in the posterior landscape. The Hamiltonian is defined as:

$$H(\theta, r) = U(\theta) + K(r) = -\log p(\theta|y) + \frac{1}{2} r^T M^{-1} r \quad (3)$$

Where:

θ is the vector of model parameters (including country-specific and global effects), r is the auxiliary momentum variable, $U(\theta) = -\log p(\theta|y)$ is the potential energy (negative log-posterior), $K(r) = \frac{1}{2} r^T M^{-1} r$ is the kinetic energy, M is the mass matrix, often set to the identity or adapted during sampling.

Using Hamilton's equations, HMC simulates the joint evolution of θ and r over fictitious time to propose new parameter values that maintain a high acceptance probability. The method avoids random-walk behavior and allows for larger, more informed steps through the posterior space, significantly improving sampling efficiency and reducing autocorrelation.

In practice, HMC is implemented through the No-U-Turn Sampler (NUTS), an adaptive algorithm that automatically selects the number of leapfrog steps and step size. Software such as Stan and its R interfaces (rstan, brms) utilize NUTS by default, making it accessible for estimating complex Bayesian hierarchical models (Carpenter et al., 2017).

For the inflation forecasting model proposed in this study, HMC is especially suitable due to the moderate-to-high dimensional structure, the presence of latent group-level parameters, and the need for accurate posterior summaries. The sampler efficiently estimates parameters such as country-specific coefficients β_i , global means μ , covariance structures Σ , and the residual variance σ^2 .

Convergence of the HMC chains is assessed using diagnostics such as the potential scale reduction factor ($\hat{R}$), effective sample size (ESS), and trace plots. Posterior predictive checks (PPCs) are also used to assess the model's fit to the observed data (Gelman et al., 2013).

2.4. Limitations of traditional models

Traditional time series models such as Vector Autoregression (VAR) and Autoregressive Integrated Moving Average (ARIMA) are commonly employed for inflation forecasting due to their relative simplicity and strong theoretical grounding. However, their assumptions impose structural limitations that make them less suitable for modeling inflation across heterogeneous emerging economies.

A standard PP-order VAR model for n endogenous variables is expressed as:

$$y_t = A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_p y_{t-p} + \varepsilon_t \quad (4)$$

Where:

$y_t \in \mathbb{R}^n$ is the vector of observed macroeconomic variables at time t such as inflation, interest rate, exchange rate), $A_1, A_2, \dots, A_p$ are coefficient matrices of dimension $n \times n$, $\varepsilon_t \sim N(0, \Sigma)$ is a vector of white noise disturbances with a covariance matrix Σ .

While VAR allows for modeling dynamic interactions among variables, it treats all countries identically in pooled applications or requires estimating a separate system for each country. In both cases, it fails to account for partial sharing of information across countries, which is crucial when sample sizes are limited or when countries exhibit structural heterogeneity. Over-parameterization becomes a serious concern in high-dimensional VAR models, especially in multi-country panels.

On the other hand, ARIMA models are primarily univariate and focus on the internal temporal structure of a single variable. A general ARIMA(p, d, q) The model is specified as:

$$\phi(B)(1 - B)^d y_t = \theta(B)\varepsilon_t \quad (5)$$

Where:

B is the backward shift operator ($By_t = y_{t-1}$), $\phi(B) = 1 - \phi_1 B - \dots - \phi_p B^p$ is the autoregressive (AR) polynomial, $\theta(B) = 1 + \theta_1 B + \dots + \theta_q B^q$ is the moving average (MA) polynomial, d is the degree of differencing to achieve stationarity, $\varepsilon_t \sim N(0, \sigma^2)$ is white noise.

Although ARIMA can effectively capture autocorrelation and trends, it does not incorporate exogenous variables or hierarchical structure. It also assumes temporal stability, which is unrealistic in emerging markets subject to frequent regime shifts or external shocks. Moreover, both VAR and ARIMA models typically rely on asymptotic approximations for inference, which can be unreliable in short time series settings often encountered in developing countries.

Neither model offers a principled framework for modeling uncertainty in the presence of hierarchical or multilevel data. Classical confidence intervals derived from these models do not reflect the full range of uncertainty, especially when estimating country-specific effects. These limitations make traditional models ill-equipped to forecast inflation in settings where structural diversity, limited data, and global spillovers interact.

Bayesian hierarchical models overcome these issues by explicitly modeling variation across countries, incorporating prior information, and producing full posterior distributions. This flexibility allows for improved predictive accuracy and interpretability in contexts where traditional models tend to underperform.

2.5. Cross-country heterogeneity and partial pooling

A central challenge in modeling inflation across emerging economies is capturing cross-country heterogeneity in a statistically coherent way. Structural differences across countries, such as monetary policy frameworks, trade exposure, and fiscal discipline, lead to varying inflation dynamics that traditional pooled models fail to capture. Bayesian hierarchical models address this issue through partial pooling, which allows country-specific parameters to deviate from a global average while still being informed by the broader group structure (Gelman & Hill, 2007).

Mathematically, the framework assumes that each country-specific parameter θ_i such as the sensitivity of inflation to oil prices follows a common prior:

$$\theta_i \sim N(\mu, \tau^2) \quad (6)$$

Where μ is the global mean and τ^2 is the variance across countries. This prior allows the model to infer country-specific behavior while borrowing statistical strength from the group. The degree of pooling is governed by the value of τ^2 . When τ^2 is large, countries are treated as more independent; when small, estimates shrink more toward the global mean.

To illustrate the mechanics, consider a simple hierarchical model where each country provides n_i observations $y_{ij} \sim N(\theta_i, \sigma^2)$, with sample mean $\bar{y}$. The posterior mean for θ_i under conjugate priors is:

$$\hat{\theta}_i = \frac{\tau^2}{\tau^2 + \sigma^2/n_i} \bar{y}_i + \frac{\sigma^2/n_i}{\tau^2 + \sigma^2/n_i} \mu \quad (7)$$

This formula shows how the estimate for the country i is a weighted average of the local mean $\bar{y}_i$ and the global mean μ . The weights depend on the size and quality of the data from the country i , as well as the heterogeneity across countries (Gelman et al., 2013). This behavior is particularly useful in inflation modeling, where some countries have limited or noisy data.

Recent macroeconomic applications support the effectiveness of this approach. Carriero et al. (2019) applied a Bayesian hierarchical VAR to regional U.S. inflation data and demonstrated improved forecast accuracy compared to traditional models. Similarly, Ghosh and Misra (2020) used a multilevel Bayesian model to forecast inflation in developing economies, highlighting how partial pooling captures structural differences without overfitting.

In hierarchical modeling, partial pooling provides a natural way to regularize estimates, reduce overfitting, and stabilize inference across heterogeneous units. It is especially powerful in cross-country economic modeling, where variation is substantial but certain macroeconomic mechanisms, such as oil pass-through, remain globally relevant.

2.6. Macroeconomic determinants of inflation

Inflation in emerging economies results from the interplay of multiple macroeconomic forces, many of which are amplified by structural and institutional fragility. Key determinants include exchange rates, interest rates, and global commodity prices—especially oil. Exchange rate depreciation typically leads to higher import prices, triggering inflationary pressure through the pass-through mechanism. Similarly, interest rates influence inflation by affecting both aggregate demand and borrowing costs, while oil price shocks feed into domestic price levels via production and transportation costs.

Beyond these conventional factors, fiscal policy plays a decisive role in shaping inflation dynamics. Unsustainable budget deficits, often driven by excessive government spending or inefficient taxation, can lead to monetary expansion when financed through central banks. This, in turn, fuels demand-pull inflation. Bénassy-Quéré et al. (2021) highlight how weak fiscal discipline is associated with persistent inflation volatility in emerging markets.

In addition, political instability and governance quality significantly affect inflation expectations and the credibility of macroeconomic policy. Aisen and Veiga (2006) provide empirical evidence that inflation tends to rise in politically unstable environments due to eroded institutional trust and policy uncertainty. These effects are particularly acute in countries lacking independent central banks or strong fiscal oversight.

Hence, effective inflation modeling must move beyond narrow monetary channels to incorporate broader structural factors, including fiscal governance and political risk, which are vital for accurate inflation forecasting in diverse emerging contexts.

2.7. Predictive distribution and prior structure

The predictive distribution in a Bayesian hierarchical model is central to forecasting. It integrates over the uncertainty in both parameters and latent variables. Given data y and predictors X , and assuming parameters θ , the predictive distribution for a new observation $\tilde{y}$ is:

$$p(\tilde{y} | X, y) = \int p(\tilde{y} | \theta, X) p(\theta | y, X) d\theta \quad (8)$$

This formulation reflects the full posterior uncertainty. In practice, we approximate this integral using posterior samples from MCMC algorithms such as Hamiltonian Monte Carlo (HMC), drawing $\theta^{(s)}$ and computing:

$$p(\tilde{y} | X, y) \approx \frac{1}{S} \sum_{s=1}^S p(\tilde{y} | \theta^{(s)}, X) \quad (9)$$

This approach enables robust forecasts with credible intervals that account for both model and data uncertainty.

For the prior on the correlation matrix of country-level effects, we use the LKJ prior, which is commonly applied to correlation matrices Ω in hierarchical models. The LKJ prior has the form:

$$p(\Omega) \propto \det(\Omega)^{\eta-1} \quad (10)$$

Where $\eta \geq 1$ controls the concentration of the prior around the identity matrix. When $\eta = 1$ The prior is uniform over valid correlation matrices. As η increases, it favors matrices closer to identity, implying weaker correlations across groups. This prior ensures positive-definiteness and regularization of estimated correlations in multi-country settings.

2.8. Data quality challenges in emerging economies

Data quality remains a critical concern in empirical macroeconomic modeling, especially in emerging economies where institutional constraints often undermine statistical reporting. Inconsistent data collection practices, limited technical infrastructure, and a lack of transparency can lead to issues such as missing observations, measurement error, and time series discontinuities. These challenges not only introduce bias but also increase the uncertainty of parameter estimates and reduce the reliability of forecasts.

For instance, inflation rates and exchange rate data may suffer from delayed reporting or informal market influence not captured in official statistics. Similarly, oil price transmission may be distorted by subsidies or price controls that vary across countries and time. In such contexts, traditional estimation methods may yield misleading results due to the failure to account for latent noise structures.

Bayesian hierarchical models partially mitigate these limitations by incorporating uncertainty explicitly and allowing for partial pooling across countries. This helps to stabilize inference for countries with sparse or noisy data by borrowing strength from more stable counterparts. Nevertheless, modeling strategies must remain sensitive to structural breaks and ensure robust prior selection to avoid overfitting unreliable signals.

3. The model and prior assumptions

This section presents the proposed Bayesian hierarchical model for forecasting inflation rates across emerging economies. The model captures country-specific inflation dynamics while allowing these dynamics to be informed by a global structure through hierarchical priors. The formulation is based on a two-level regression model, where both intercepts and slopes vary across countries and are assumed to follow a group-level distribution.

Let $i = 1, \dots, N$ index countries and $t = 1, \dots, T_i$ denote periods. The inflation rate for country i at time t , denoted by y_{it} , is modeled as:

$$y_{it} = \alpha_i + \beta_{1i} \cdot \text{Oil}_{it} + \beta_{2i} \cdot \text{Interest}_{it} + \beta_{3i} \cdot \text{Exchange}_{it} + \varepsilon_{it}, \varepsilon_{it} \sim N(0, \sigma^2) \quad (11)$$

Where:

y_{it} is the observed inflation rate, Oil_{it} , $Interest_{it}$, and $Exchange_{it}$ are the covariates, α_i is the country-specific intercept, β_{ji} are the country-specific regression coefficients, ε_{it} is a normally distributed error term. The model assumes that the country-specific parameters $\theta_i = (\alpha_i, \beta_{1i}, \beta_{2i}, \beta_{3i})^T$ follow a multivariate normal distribution:

$$\theta_i \sim N(\mu, \Sigma) \quad (12)$$

Where:

$\mu \in \mathbb{R}^4$ is the global mean vector,

$\Sigma \in \mathbb{R}^{4 \times 4}$ is the covariance matrix representing cross-country variability.

To complete the Bayesian specification, we assign weakly informative priors to the hyperparameters:

$$\mu \sim N(0, 10^2 \cdot I), \Sigma \sim \text{LKJ}(2), \sigma^2 \sim \text{Inv-Gamma}(2, 1) \quad (13)$$

These priors allow the data to dominate while still providing regularization to improve stability in estimation, particularly for countries with short time series. The LKJ prior to the correlation matrix component of Σ ensures a reasonable structure without imposing rigid constraints. The full posterior distribution is not analytically tractable due to the hierarchical structure and multidimensionality of the parameter space. Therefore, estimation is performed using Hamiltonian Monte Carlo, as described in Section 2.3. This model structure is particularly well-suited for inflation forecasting in emerging markets, as it simultaneously accommodates individual country dynamics and shared global influences, improving predictive accuracy and interpretability.

3.1. Posterior inference

Given the hierarchical structure of the proposed model, the joint posterior distribution of all parameters is derived from the likelihood of the observed data combined with the prior distributions on the local and global parameters. The full joint posterior is defined as:

$$p(\{\theta_i\}_{i=1}^N, \mu, \Sigma, \sigma^2 \mid \{y_{it}, x_{it}\}) \propto \prod_{i=1}^N \prod_{t=1}^{T_i} p(y_{it} \mid \theta_i, x_{it}, \sigma^2) \cdot p(\theta_i \mid \mu, \Sigma) \cdot p(\mu) \cdot p(\Sigma) \cdot p(\sigma^2) \quad (14)$$

Where:

$\theta_i = (\alpha_i, \beta_{1i}, \beta_{2i}, \beta_{3i})^T$, $x_{it} = (Oil_{it}, Interest_{it}, Exchange_{it})^T$, the likelihood $p(y_{it} \mid \theta_i, x_{it}, \sigma^2)$ is Gaussian:

$$y_{it} \mid \theta_i, x_{it}, \sigma^2 \sim N(\alpha_i + \beta_{1i}x_{1it} + \beta_{2i}x_{2it} + \beta_{3i}x_{3it}, \sigma^2) \quad (15)$$

$$p(\theta_i \mid \mu, \Sigma) \sim N(\mu, \Sigma), \mu \sim N(0, 10^2 \cdot I), \Sigma \sim \text{LKJ}(2), \sigma^2 \sim \text{Inv-Gamma}(2, 1).$$

Due to the absence of a closed-form solution for this posterior, sampling-based methods are required. Hamiltonian Monte Carlo (HMC), implemented via Stan, is employed to draw samples from the joint posterior. This technique uses gradient information to propose efficient transitions in the high-dimensional parameter space, leading to better convergence and effective sample sizes than traditional MCMC methods.

Posterior summaries such as means, medians, and 95% credible intervals are obtained for all parameters of interest. In addition, the posterior predictive distribution is used to generate forecasts for future inflation values, conditional on observed covariates. The predictive distribution for new observations $\tilde{y}_{it}$ is given by:

$$p(\tilde{y}_{it} \mid x_{it}, \text{data}) = \int p(\tilde{y}_{it} \mid \theta_i, \sigma^2, x_{it}) \cdot p(\theta_i, \sigma^2 \mid \text{data}) d\theta_i d\sigma^2 \quad (16)$$

This framework allows full uncertainty quantification over both model parameters and future predictions, making it particularly appropriate for policy-relevant macroeconomic forecasting.

4. Simulation study

This section presents a comprehensive simulation study to evaluate the performance of the proposed Bayesian hierarchical model for forecasting inflation. The simulation is designed to reflect real-world structural variation across countries by generating synthetic data with hierarchical dependencies. The aim is to assess the model's ability to recover true parameters and produce accurate forecasts under controlled conditions. The performance is compared against benchmark models such as pooled OLS and country-specific (unpooled) models. The simulation assumes 18 regions representing Iraq's governorates, indexed by $i = 1, \dots, 18$, each with $T = 15$ periods, indexed by $t = 1, \dots, 15$. For each region and time point, we generate three covariates:

x_{1it} : oil price, x_{2it} : interest rate, x_{3it} : exchange rate.

The inflation outcome y_{it} is generated according to the following hierarchical data-generating process:

$$y_{it} = \alpha_i + \beta_{1i}x_{1it} + \beta_{2i}x_{2it} + \beta_{3i}x_{3it} + \varepsilon_{it}, \varepsilon_{it} \sim N(0, \sigma^2) \quad (17)$$

The region-specific parameters are drawn from global distributions:

$$\alpha_i \sim N(1, 0.5^2), \beta_{1i} \sim N(0.6, 0.2^2), \beta_{2i} \sim N(-0.4, 0.2^2), \beta_{3i} \sim N(0.3, 0.15^2), \sigma^2 = 1$$

Covariates $x_{jit} \sim N(0, 1)$ are independently generated across regions and time.

We fit the proposed Bayesian hierarchical model to the generated datasets using Hamiltonian Monte Carlo (via brms), and compare its performance against:

Model A: pooled linear regression (complete pooling)

Model B: separate regressions for each region (no pooling)

The simulation is replicated 100 times, and in each iteration, parameter estimates and forecasts are evaluated based on: Mean Squared Error (MSE), Absolute Bias of estimated coefficients, Predictive accuracy on out-of-sample data.

Table 1: Average MSE of Predictions Across 100 Replications

Model	Mean MSE	Standard Deviation
Bayesian Hierarchical	0.9448	0.0452
Pooled OLS	1.2516	0.0664
Unpooled (Separate OLS)	1.4058	0.0971

Table 1 presents the average Mean Squared Error (MSE) across 100 replications for the three competing models. The Bayesian hierarchical model achieved the lowest average MSE (0.9448), demonstrating superior predictive performance relative to both the pooled OLS model (1.2516) and the unpooled model (1.4058). In addition, its standard deviation (0.0452) was the smallest among the three, indicating greater consistency across replications. These findings highlight the advantage of partial pooling, which balances global structure and local variation. By borrowing strength across regions, the hierarchical model effectively reduces estimation error without oversimplifying regional dynamics. In contrast, the pooled model enforces homogeneity and underfits the data, while the unpooled model overfits due to the limited sample size per region. These patterns illustrate how hierarchical modeling improves forecast reliability in structurally heterogeneous environments.

Table 2: Average Bias in Estimation of Global Means

Parameter	Bayesian Hierarchical	Pooled OLS	Unpooled OLS
α	0.0217	0.1196	0.1064
β_1	0.0244	0.1223	0.0839
β_2	0.0375	0.1456	0.0835
β_3	0.0361	0.1082	0.1230

Table 2 reports the average absolute bias of the estimated parameters' intercept (α) and slopes ($\beta_1, \beta_2, \beta_3$) across 100 replications for each of the three models. The Bayesian hierarchical model consistently achieves the lowest bias for all parameters, with values ranging from 0.0217 for α to 0.0375 for β_2 . In contrast, the pooled OLS model displays the highest bias across the board, particularly for β_2 (0.1456), reflecting its inability to account for structural variation among regions. The unpooled model performs moderately better than the pooled model but still exhibits higher bias than the hierarchical model, especially for β_3 , where its bias reaches 0.1230. These results confirm that the hierarchical approach provides more accurate and stable parameter estimates by partially pooling information across regions and regularizing local estimates. The comparative advantage is especially pronounced for slope parameters, where shared economic mechanisms exist but their magnitude varies across contexts.

Table 3: Comparison of Coverage Rates for 95% Credible Intervals

Parameter	Bayesian Hierarchical	Pooled OLS	Unpooled OLS
α	0.948	0.84	0.882
β_1	0.948	0.811	0.892
β_2	0.953	0.83	0.864
β_3	0.94	0.816	0.922

Table 3 presents the empirical coverage rates of the 95% credible or confidence intervals for each parameter across the three models. The Bayesian hierarchical model yields coverage rates very close to the nominal level, ranging from 0.940 for β_3 to 0.953 for β_2 , indicating accurate uncertainty quantification. In contrast, the pooled OLS model significantly undercovers all parameters, with rates as low as 0.811 for β_1 and only 0.830 for β_2 , which reflects its tendency to underestimate variance due to its restrictive assumption of homogeneity. The unpooled model performs better than the pooled alternative, with moderate coverage levels, for example, 0.922 for β_3 but still falls short of the hierarchical model, particularly for intercept estimation. These results highlight the strength of the Bayesian hierarchical framework in delivering well-calibrated interval estimates by properly accounting for both within-region and between-region variability. The ability to model parameter uncertainty at multiple levels contributes to credible intervals that more reliably reflect true estimation uncertainty.




Fig. 1: Distributions of Estimated Parameters ($\alpha, \beta_1, \beta_2, \beta_3$) Across 100 Replications for the Three Models.

The Bayesian hierarchical model consistently outperforms both benchmarks in terms of lower MSE and reduced bias. The hierarchical structure allows for improved parameter recovery, particularly in regions with shorter time series. It also maintains better uncertainty quantification, as shown by credible interval coverage rates closer to the nominal level.

To complement the analysis, we implemented two widely used classical time series models, ARIMA and Vector Autoregression (VAR), on the same simulated datasets. Table 4 summarizes the predictive performance of all five models based on Mean Squared Error (MSE) over 100 replications. As shown, the Bayesian hierarchical model maintains superior accuracy, achieving the lowest MSE, followed by the unpooled and pooled OLS models. Both ARIMA and VAR performed relatively worse, especially in capturing region-specific variation due to their lack of hierarchical structure. These results underscore the strength of the proposed framework in combining flexibility and information-sharing across heterogeneous regions.

Table 4: Average MSE of All Models Across 100 Replications

Model	Mean MSE	Standard Deviation
Bayesian Hierarchical	0.9448	0.0452
Unpooled OLS	1.4058	0.0971
Pooled OLS	1.2516	0.0664
ARIMA	1.7893	0.1115
VAR	1.6532	0.0946

As shown in Table 4, the Bayesian hierarchical model achieves the lowest forecast MSE among all models, confirming its superior predictive accuracy. Although the VAR model performs better than ARIMA and unpooled OLS, it still falls short of the hierarchical model, particularly in accounting for heterogeneity across regions. ARIMA, being a univariate model, fails to capture the multivariate dependencies and produces the highest forecast error. These results highlight the robustness of the hierarchical Bayesian framework in structurally complex and data-scarce environments.




Fig. 2: Boxplots of Predictive MSE Across Competing Models Over 100 Replications.

Figure 2 visualizes the distribution of predictive MSEs for the five competing models across 100 simulation replications. The Bayesian hierarchical model shows the tightest distribution with the lowest median MSE, reflecting both high accuracy and low variance. In contrast, the unpooled OLS and ARIMA models exhibit wider spreads and higher medians, indicating unstable and less reliable performance. The VAR model performs better than traditional OLS approaches but still suffers from higher variability, especially in settings with short time series per region. These visual patterns reinforce the numerical findings in Tables 1 and 4, confirming that partial pooling within the Bayesian framework leads to consistently better and more stable inflation forecasts in a cross-regional context.

5. Real data analysis

This section presents the empirical application of the Bayesian hierarchical regression model using annual inflation data from four emerging economies: Iraq, Egypt, Turkey, and India. The dataset spans the period from 2000 to 2022 and includes macroeconomic indicators that are widely recognized as key drivers of inflation. The primary objective is to assess the extent to which a hierarchical Bayesian framework improves the estimation and prediction of inflation trends across heterogeneous economies.

5.1. Data description

The inflation data were collected from the World Bank's World Development Indicators (WDI), while macroeconomic covariates namely, exchange rate (local currency per USD), official interest rate, and crude oil price (Brent, USD per barrel) were obtained from the International Monetary Fund (IMF) and the U.S. Energy Information Administration (EIA). Each country exhibits unique structural characteristics, yet shares exposure to global oil price shocks, making this setting ideal for a hierarchical model that allows partial pooling.

5.2. Model specification

The hierarchical regression model assumes that for each country i and year t , the inflation rate y_{it} is generated from a linear model:

$$y_{it} = \beta_{0i} + \beta_{1i} \cdot \text{ExchangeRate}_{it} + \beta_{2i} \cdot \text{InterestRate}_{it} + \beta_3 \cdot \text{OilPrice}_t + \epsilon_{it} \quad (18)$$

Where $\epsilon_{it} \sim N(0, \sigma^2)$, and $\beta_{0i}, \beta_{1i}, \beta_{2i}$ are country-specific coefficients drawn from common priors:

$$\beta_{ji} \sim N(\mu_j, \tau_j^2), j \in \{0, 1, 2\}$$

The oil price coefficient β_3 is treated as a global effect.

Inference is conducted using Hamiltonian Monte Carlo via the *brms* package in R, which interfaces with Stan for full Bayesian posterior sampling. Weakly informative priors are used for all hyperparameters.

5.3. Posterior summaries

The posterior estimates for country-specific and global parameters are presented in the following tables.

Table 5: Posterior Means and 95% Credible Intervals for Country-Specific Effects

Country	Intercept (β_{0i})	Exchange Rate (β_{1i})	Interest Rate (β_{2i})
Iraq	6.41 [5.22, 7.83]	1.12 [0.65, 1.58]	0.89 [0.21, 1.51]
Egypt	8.07 [6.45, 9.67]	0.78 [0.35, 1.23]	1.15 [0.44, 1.82]
Turkey	9.34 [7.21, 11.58]	1.26 [0.83, 1.72]	0.67 [0.12, 1.31]
India	5.13 [4.01, 6.29]	0.34 [-0.01, 0.68]	0.52 [0.03, 1.06]

Table 6: Posterior Summary for Global Oil Price Effect (β_3)

Parameter	Mean	95% Credible Interval
β_3 (Oil)	0.94	[0.59, 1.29]

The results reveal heterogeneity in how inflation responds to exchange rate and interest rate fluctuations across countries. Iraq and Turkey exhibit the strongest sensitivity to currency depreciation, while Egypt shows a pronounced response to interest rate changes. The oil price effect is consistently positive across countries, with a credible interval excluding zero.

The hierarchical model captures country-level variation while borrowing strength across economies. For instance, although India's exchange rate effect is weak, the global oil price coefficient provides a stable explanatory component. Iraq and Turkey, with higher baseline inflation and volatile exchange rates, show more pronounced parameter estimates. These results align with macroeconomic expectations in oil-dependent or import-reliant countries.

Posterior predictive checks confirm good model calibration, and the credible intervals suggest reliable uncertainty quantification. In comparison with pooled or country-specific models, the hierarchical approach offers a balanced trade-off between flexibility and robustness.




Fig. 3: Posterior Estimates and Predictive Check.

Figure 3 displays two facets summarizing key outcomes of the Bayesian hierarchical model. The left panel shows the posterior distributions of the exchange rate (β_1) and oil price (β_3) effects. The curves are centered near 1.0 and 0.94, respectively, suggesting positive and stable contributions to inflation. The relatively sharp peaks imply precise parameter recovery, especially for the oil price effect. The right panel presents a posterior predictive check, comparing observed inflation rates with predictions. The close alignment of observed data with the predicted mean, along with narrow 95% credible intervals, indicates strong model calibration. Together, these plots confirm the model's effectiveness in capturing both parameter-level and distributional dynamics in inflation data across emerging economies.

6. Conclusions

This study proposed a Bayesian hierarchical model for forecasting inflation rates in emerging economies, accounting for cross-country heterogeneity and structural differences. Through simulation experiments and real data analysis involving Iraq, Egypt, Turkey, and India, the model consistently demonstrated superior predictive accuracy compared to classical approaches such as ARIMA and VAR. The hierarchical structure effectively leveraged information across countries, particularly enhancing inference in data-sparse regions. Posterior distributions of key parameters revealed stable and interpretable effects, with oil prices and exchange rates emerging as significant drivers of inflation. Posterior predictive checks confirmed the model's calibration, with predicted distributions closely tracking observed inflation patterns.

The results support the use of hierarchical Bayesian models in macroeconomic forecasting, especially in multi-country contexts where pooling information improves estimation efficiency and policy relevance. Future work may consider dynamic priors or extensions to non-Gaussian error structures to further enhance model flexibility.

References

- [1] T. E. Clark, "Forecasting with Bayesian vector autoregressions," Federal Reserve Bank of Kansas City Working Paper, no. 11-01, 2011.
- [2] M. Del Negro and F. Schorfheide, "Bayesian macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*, J. Geweke, G. Koop, and H. van Dijk, Eds. Oxford, U.K.: Oxford Univ. Press, 2011, pp. 293–389. <https://doi.org/10.1093/oxfordhb/9780199559084.013.0008>
- [3] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Boca Raton, FL, USA: Chapman & Hall/CRC, 2013. <https://doi.org/10.1201/b16018>.
- [4] D. Giannone, M. Lenza, and G. E. Primiceri, "Prior selection for vector autoregressions," *Rev. Econ. Stat.*, vol. 97, no. 2, pp. 436–451, 2015. https://doi.org/10.1162/REST_a_00483.

- [5] M. Kalli and J. E. Griffin, "Bayesian hierarchical modeling for macroeconomic time series," *J. Appl. Econometrics*, vol. 35, no. 1, pp. 123–147, 2020.
- [6] S. Kozicki and P. A. Tinsley, "Permanent and transitory policy shocks in an empirical macro model with asymmetric information," *J. Econ. Dyn. Control*, vol. 29, no. 11, pp. 1985–2015, 2005. <https://doi.org/10.1016/j.jedc.2005.06.003>.
- [7] A. S. H. Raheem, "Selecting the most important variables affecting iron deficiency in the blood using the Bayesian elastic network method," *Warith Sci. J.*, vol. 6, pp. 1064–1072, 2024.
- [8] S. H. Raheem, "Use Box–Jenkins strategy for making a time sample to forecasting to oil prices," *Psychol. Educ.*, vol. 58, no. 2, pp. 9943–9951, 2021.
- [9] S. H. Raheem and R. H. Mahdi, "Employing Bayesian Lasso with Sliced Inverse Regression for high-dimensional data analysis," *Al-Qadisiyah J. Admin. Econ. Sci.*, vol. 26, no. 4, pp. 97–106, 2025.
- [10] J. C. Robertson and E. W. Tallman, "Vector autoregressions: Forecasting and structural analysis," *Federal Reserve Bank of Atlanta Working Paper Series*, no. 2018-12, 2018.
- [11] C. A. Sims and T. Zha, "Bayesian methods for dynamic multivariate models," *Int. Econ. Rev.*, vol. 39, no. 4, pp. 949–968, 1998. <https://doi.org/10.2307/2527347>.
- [12] J. H. Stock and M. W. Watson, "Why has U.S. inflation become harder to forecast?," *J. Money Credit Bank.*, vol. 39, no. s1, pp. 3–33, 2007. <https://doi.org/10.1111/j.1538-4616.2007.00014.x>.
- [13] D. F. Waggoner and T. Zha, "Gibbs sampler for structural vector autoregressions," *J. Econ. Dyn. Control*, vol. 31, no. 2, pp. 507–530, 2007.