Journal of Economic Theory and Econometrics, Vol. 36, No. 3, Sep. 2025, 61–105

Bayesian Forecasting of China's Housing Prices: Dynamic Factor Identification and Predictive Evaluation*

Xia Gao†

Abstract This paper applies a Bayesian Variable Selection (BVS) framework to forecast the year-on-year growth rate of China's newly built housing price index (HPI YoY). Using a broad set of macroeconomic and financial predictors, we implement a hierarchical BVS model with rolling-window estimation and direct multi-horizon forecasting. Out-of-sample performance is evaluated against autoregressive (AR) models, the random walk (RW), and a machine learning benchmark, the Random Forest (RF). The results show that BVS consistently outperforms AR and RW across most horizons in both point and density forecasts, and dominates RF within two years ($h = 1$–18), while RF is only slightly better at very long horizons ($h = 24, 30, 36$). Horizon-specific predictors provide further insights: lagged HPI and market sentiment (RECI) drive short-run dynamics; RECI, inflation (CPI), and housing credit conditions (HPF, IHLL) matter at medium to long horizons; and demographic fundamentals (PNGR) dominate at long and ultra-long horizons. The forecast results point to a subdued housing market over the next two years, consistent with a structural break around April 2022 and the lasting impact of demographic shifts. Overall, the study demonstrates that BVS not only improves forecasting accuracy but also enhances interpretability, making it a valuable tool for academic research and housing market policy design.

Keywords Bayesian variable selection, housing price forecast, macroeconomic predictors, out-of-sample evaluation, China.

JEL Classification C11, C53, R31.

*I am deeply grateful to the editor and two anonymous reviewers for their thoughtful comments and suggestions. I also thank Kyu Ho Kang and Jungmo Yoon for their helpful feedback during the development of this paper.

†Department of Economics, Korea University, 145 Anam-ro, Seongbuk-gu, Seoul, Republic of Korea 02841. E-mail: gaofanbo@korea.ac.kr.

Received August 6, 2025, Revised September 1, 2025, Accepted September 10, 2025

1. INTRODUCTION

House prices are an important indicator of the health and stability of the real estate market. Predicting changes in house prices helps to clearly understand the real estate market. Therefore, real estate investors, real estate developers and government regulators always pay close attention to house price trends. Accurate house price forecasts not only help the government to regulate the real estate market more effectively, but also help real estate developers make smarter investment decisions. Real estate investors' decisions also largely depend on the forecast of future house prices. Especially in this period, after more than 30 years of continuous increases in Chinese house prices, they have begun to decline in the past three years. How government regulators should better stabilize house prices has become an important topic.

At the same time, compared with major developed economies such as the United States and the United Kingdom, China's real estate market is quite unique. Unlike other commodities, housing is the most important asset of a typical Chinese family. In Chinese culture, it is a key indicator of economic success and social status. Therefore, it is not only indispensable, but also often 'the more the better'. According to Huang, Yi and Clark (2020), the house ownership rate in China has exceeded 80%, compared to just over 60% in major 'ownership countries'. Moreover, more than 20% of urban households in China own multiple homes—significantly higher than the 13% in the United States and around 10% in the United Kingdom.

Underdeveloped financial markets and capital controls limit household investment options and make residential properties more like a financial product rather than a pure place to live. Therefore, when housing prices in China no longer continue to rise and begin to show a downward trend, it becomes critically important for homeowners or investment-oriented buyers to judge the trajectory of future housing prices and to understand which factors are driving these changes.

Given these unique institutional features, understanding housing price dynamics in China requires a macroeconomically grounded approach. This study develops a forecasting model at the national level, using a wide range of macro-financial variables and a Bayesian Variable Selection (BVS) framework to capture short-, medium-, and long-term drivers of housing price growth. The findings offer insights not only for central policymakers, but also provide useful guidance for real estate investors seeking to understand nationwide market trends. In addition, the proposed modeling framework lays the groundwork for future studies aiming to analyze local housing markets using the same infrastructure.

Many scholars have conducted extensive research on house price forecasting. For example, Crawford and Fratantoni (2003) forecasted quarterly house prices in five U.S. states from 1979 to 2001 using ARIMA, GARCH, and regime-switching models, and compared their forecasting performance. Their study reveals that regime-switching models are a compelling choice for real estate markets with historical boom-and-bust cycles. However, simple ARIMA models generally perform better in out-of-sample forecasting. Rapach and Strauss (2007) used an autoregressive distributed lag (ARDL) model framework with 25 determinants to forecast real housing price growth for individual states in the Federal Reserve's Eighth District. They find that ARDL models tend to outperform a benchmark AR model.

However, the traditional housing price prediction model faces at least one or two major problems. First, the traditional housing price prediction model cannot take into account the impact of multiple variables from various fields on housing prices. Second, if too many irrelevant variables are used, it will inevitably cause overfitting problems, which will affect the accuracy of the prediction. Moreover, the impact of various variables on housing prices is dynamic, and it is also necessary to dynamically reflect the effects of various factors on housing prices.

With the rise of computational tools, many scholars have also explored machine learning methods for house price prediction. For example, Adetunji et al. (2022) and Xu and Zhang (2021) apply random forest and neural networks to estimate housing values using U.S. and Chinese data, respectively, reporting high predictive accuracy. Other studies such as Satish et al. (2019), Kuvalekar et al. (2020), and Teoh et al. (2023) employ models like Lasso regression, decision trees, and explainable AI techniques to improve prediction performance.

The rapid advancement of machine learning techniques has improved the ability to forecast housing prices while mitigating overfitting risks. However, most existing machine learning studies focus on micro-level features of individual properties—such as location, structure, and amenities—rather than macroeconomic dynamics. In contrast, this study aims to forecast the overall housing price trend in China by applying a Bayesian Variable Selection (BVS) model to a broad set of macroeconomic and financial indicators. Given the increasing uncertainty in China's housing market, understanding the macroeconomic drivers of housing price fluctuations is critical for both policymakers and individual investors.

The BVS approach offers several key advantages. It dynamically selects relevant predictors across short-, medium-, and long-term horizons and captures the evolving transmission of macroeconomic influences. Unlike conventional machine learning methods that rely on fixed variable sets or produce only point forecasts, BVS generates full predictive distributions and quantifies model uncertainty through posterior inclusion probabilities. This enables transparent interpretation of variable importance and enhances decision-making under uncertainty. By applying this framework to China's housing price data, our study provides both methodological contributions and policy-relevant insights.

Kang (2018) proposed a mortgage prediction method and performed density forecasting using a Bayesian machine learning algorithm based on conditional Bayesian model averaging. Drawing inspiration from such applications of Bayesian Variable Selection (BVS) models in macroeconomic contexts, our study extends this framework to forecast the year-on-year growth rate of China's housing price index. In contrast to Kang's study, which employed an indirect prediction strategy for rolling forecasts, we adopt a direct multi-horizon forecasting approach to better capture the dynamic influence of macroeconomic variables over short-, medium-, and long-term horizons.

While Bayesian Variable Selection (BVS) methods have been widely applied in macroeconomic forecasting, their use in housing price prediction—particularly in identifying and interpreting macroeconomic influence factors—remains relatively underexplored. Prior research across various countries shows that housing prices are shaped by a wide array of factors, including GDP growth, demographics, credit and money supply, income, interest rates, inflation, speculative capital flows, taxation policies, and stock market wealth Mikhed and Zemcik (2009); Rapach and Strauss (2009); Shiller (2005); Tsatsaronis and Zhu (2004). Fundamentally, housing prices reflect supply and demand dynamics: higher interest rates raise borrowing costs for both developers and households, thereby constraining investment and demand Arcelus and Meltzer (1973); Smith (1969); Maisel (1968). Over the long run, prices tend to move with macroeconomic development, as economic growth, urbanization, and rising incomes expand housing demand Gottlieb (1976); Liang and Gao (2007). In China, monetary policy—through interest rate and deposit regulation—exerts particularly strong influence on housing prices Cai, Chen and Wang (2023), motivating the selection of macro-financial variables to forecast changes in the national housing price index.

The purpose of this study is to analyze the key macroeconomic variables that dynamically influence changes in housing prices using Bayesian Variable Selection techniques, and to compare and evaluate the models' predictive performance through out-of-sample forecasts. We use the year-on-year growth rates of China's New Housing Price Index as the dependent variables in our forecasting models.

In our implementation of the Bayesian Variable Selection (BVS) technique, we build on the original approach proposed by George and McCulloch (1993), and adopt a common extension in which the hyperparameters b_0 and b_1 are treated as unknown and estimated from the data. This hierarchical prior setting helps mitigate the sensitivity of results to fixed hyperparameter choices and enhances the model's adaptability to changing macroeconomic environments Brown, Vannucci and Fearn (1998). To further improve out-of-sample forecasting performance and account for model uncertainty, we propose two alternative specifications—BVS1 and BVS2—each based on distinct prior configurations. Our empirical results confirm that the BVS approach effectively captures horizon-dependent macroeconomic drivers and outperforms benchmark models across all forecast horizons.

The structure of the rest of the paper is as follows. Section 2 outlines the methodology, presenting the Bayesian Variable Selection (BVS) model and the Gibbs sampling algorithm. Section 3 describes the predictive distribution and the out-of-sample forecasting framework. Section 4 introduces the dataset and provides details on the predictor variables. Section 5 presents the empirical analysis, dynamically examining the influencing factors and illustrating the out-of-sample forecast performance. Section 6 provides a robustness check by detecting structural breaks in the housing price series using a Bayesian change point model. Finally, Section 7 concludes the paper and discusses broader policy implications.

2. BAYESIAN VARIABLE SELECTION MODEL

In a macroeconomic forecasting, there are often too many potentially relevant predictors. For instance, most macroeconomic variables are theoretically related to Exchange rates in a general equilibrium context. Likewise, when trying to predict housing prices, there are too many variables to consider, such as the current housing price, the economy, household income, inflation rate, loan interest rates, employment rate, and population changes. Among a number of regressors, only a few variables have a signicant prediction power, while the others are unnecessary. But inclusion of unnecessary regressors causes a poor prediction performance.

In practice, it is difficult to select important regressors maximizing the predictive accuracy. Because if the number of regressors is 10, then the number of models to be compared is $2^{10} = 1024$. But in most cases of macroeconomic forecasting, it is larger than 10. It needs too many models to compare.

Bayesian(or stochastic) variable selection can solve the above problem well by selecting relevant predictors among a number of regressors. This is a statistical learning algorithm designed to automatically select important regressors by the information in data. The key idea is to shrink irrelevant regression coe cients towards zero.

2.1. HIERARCHICAL BAYESIAN VARIABLE SELECTION MODEL

To obtain robust variable selection in high-dimensional macroeconomic forecasting, we adopt a hierarchical spike-and-slab BVS specification. Conventional implementations fix the hyperparameters of the spike-and-slab prior—the prior inclusion probability p and the spike/slab variances (b_0, b_1) —which makes results sensitive to tuning and can lead to over- or under-selection of predictors, thereby hurting predictive performance.[1] To mitigate this sensitivity, we endow p, b_0, and b_1 with hyperpriors and estimate them jointly with the regression parameters, following George and McCulloch (1997) and Ishwaran and Rao (2005). This allows the degree of sparsity and shrinkage to adapt to the signal in the data and improves out-of-sample stability.

Let K denote the number of predictors, and let β_k denote the coefficient on predictor k $(k = 1, \ldots, K)$. The hierarchical prior is specified as follows:

$$p \sim \text{Beta}(a_0, c_0),$$

$$Pr[\gamma_k = 1 \mid p] = p \text{ and } Pr[\gamma_k = 0 \mid p] = 1 - p.$$

[1]See O'Hara and Sillanpää (2009) for a review of the sensitivity of BVS methods to prior settings.

$$
\begin{aligned}
b_0 &\sim \text{InverseGamma}(\alpha_{00}/2, \delta_{00}/2), \\
b_1 &\sim \text{InverseGamma}(\alpha_{01}/2, \delta_{01}/2), \\
b_k &= \gamma_k \times b_1 + [1-\gamma_k] \times b_0, \\
\beta_k \mid \gamma_k, b_0, b_1 &\sim \mathcal{N}(\beta_0 = 0,\ b_k) \ \text{ for } k = 1, 2, \ldots, K.
\end{aligned}
$$

Under this construction, when $\gamma_k = 1$, the prior variance $b_k = b_1$ is large, allowing β_k to deviate significantly from zero—this corresponds to the "slab" component. Conversely, when $\gamma_k = 0$, the prior variance $b_k = b_0 \approx 0$, heavily shrinking β_k toward zero—this is the "spike" component. This flexible prior specification is widely referred to as the spike-and-slab prior.[2]

Once $\{\beta_k\}_{k=1}^K$ are generated, the dependent variable Y is assumed to follow a conditional normal distribution as in the preliminary model:

$$
\sigma^2 \sim \text{InverseGamma}\left(\frac{\alpha_0}{2}, \frac{\delta_0}{2}\right),
$$

$$
Y \mid X, \beta, \sigma^2 \sim \mathcal{N}(X\beta, \sigma^2 I_T).
$$

2.2. PARAMETER POSTERIOR SAMPLING FOR HIERARCHICAL BVS MODEL

The model is a hierarchical model, but since it is based on a linear regression model, it does not require a complex algorithm and has the advantage of being able to be simply estimated using the Gibbs sampling technique. Gibbs sampling samples parameters from the full conditional distribution of each parameter in each iteration of the simulation. The algorithm are summarized as Table 1.

3. DISTRIBUTION PREDICTION USING BVS MODEL

3.1. PREDICTION MODEL STRUCTURE

We begin by formulating the Bayesian Variable Selection (BVS) model used to generate multi-horizon forecasts of housing price growth. The predictive distribution for the h-step-ahead forecast is defined as:

[2]The term "spike-and-slab" was first introduced by Mitchell and Beauchamp (1988). See also Ishwaran and Rao (2005) for a detailed treatment.

Algorithm: Bayesian Variable Selection

Step 0	Initialize the parameters and set $j = 1$.
Step 1	Sample $\beta^{(j)}$ from $\beta \mid Y, \sigma^{(j-1)}, b_1^{(j-1)}, b_0^{(j-1)}, \Gamma^{(j-1)} \sim \mathcal{N}(\bar{\beta}, \bar{B})$.
Step 2	Sample $\sigma^{(j)}$ from $\sigma^2 \mid Y, \beta^{(j)} \sim \text{InverseGamma}\left(\frac{\alpha_0+T}{2}, \frac{\delta_0+\left(Y-X\beta^{(j)}\right)'\left(Y-X\beta^{(j)}\right)}{2}\right)$.
Step 3	Sample $b_0^{(j)}$ and $b_1^{(j)}$ from $b_0 \mid \beta^{(j)}, \Gamma^{(j-1)} \sim \text{InverseGamma}\left(\frac{\alpha_{00}+K_0}{2}, \frac{\delta_{00}+\beta^{(j)\prime}_{\Gamma^{(j-1)}=0}\beta^{(j)}_{\Gamma^{(j-1)}=0}}{2}\right)$, $b_1 \mid \beta^{(j)}, \Gamma^{(j-1)} \sim \text{InverseGamma}\left(\frac{\alpha_{01}+K_1}{2}, \frac{\delta_{01}+\beta^{(j)\prime}_{\Gamma^{(j-1)}=1}\beta^{(j)}_{\Gamma^{(j-1)}=1}}{2}\right)$.
Step 4	Sample $p^{(j)}$ from $p \mid \Gamma \sim \text{Beta}(a_0 + K_1, c_0 + K_0)$.
Step 5	For each k, if $u_k \sim \text{Unif}(0,1)$ is less than $$\frac{p^{(j)}\mathcal{N}\left(\beta_k^{(j)} \mid 0, b_1^{(j)}\right)}{p^{(j)}\mathcal{N}\left(\beta_k^{(j)} \mid 0, b_1^{(j)}\right) + \left(1-p^{(j)}\right)\mathcal{N}\left(\beta_k^{(j)} \mid 0, b_0^{(j)}\right)},$$ then set $\gamma_k^{(j)} = 1$, and 0 otherwise.
Step 6	Repeat Steps 1–5 n times.

Table 1: Algorithm steps for Bayesian Variable Selection. This table summarizes the Gibbs sampling procedure for the hierarchical BVS model.

$$y_{t+h} \mid \mathcal{F}_t, \theta \sim \mathcal{N}\left(\beta_0 + \sum_{k=1}^{K} X_{k,t}\beta_k \cdot \gamma_k,\ \sigma^2\right).$$

Here, $\mathcal{F}_t$ denotes the information set available up to time t, The parameter vector θ includes all model parameters. β_0 represents the intercept term. $X_{k,t}$ denotes the kth lag-1 predictor, β_k is the coefficient for the k-th predictor and γ_k is the binary inclusion indicator that determines whether the kth variable is selected into the model. σ^2 is the variance term.

The predictive distribution is therefore a normal distribution centered at the linear combination of selected predictors and the intercept, with variance σ^2. These parameters are estimated through Gibbs sampling under the BVS algorithm, and the resulting posterior distributions characterize the full predictive uncertainty.

To improve forecast accuracy, we extend the model to include multiple lags for each predictor, up to a maximum lag order J. The extended prediction model is:

$$y_{t+h} \mid \mathcal{F}_t, \theta \sim \mathcal{N}\left(\beta_0 + \sum_{j=1}^{J}\sum_{k=1}^{K} X_{k,t-j+1}\beta_{j,k}\cdot\gamma_{j,k},\ \sigma^2\right),$$

where β_0 is the intercept term; $X_{k,t-j+1}$ is the value of the k-th predictor at lag j; $\beta_{j,k}$ is the coefficient associated with the k-th predictor at lag j; $\gamma_{j,k}$ is the inclusion indicator for the k-th predictor at lag j; and σ^2 is the posterior variance of the error term.

This formulation enables the model to flexibly adapt to time-lagged dynamics, while the Bayesian Variable Selection algorithm ensures a parsimonious representation by including only the most informative predictors. To implement this structure, we horizontally concatenate lagged values of each predictor from $j = 1$ to J into the predictor matrix X_t.[3]

The time index t refers to the forecast origin during the out-of-sample evaluation window. In our study, the out-of-sample period consists of 36 months, from May 2022 to April 2025. We compare forecasting performance across different lag orders J to identify the optimal dynamic structure for each horizon.

3.2. OUT-OF-SAMPLE FORECASTING FRAMEWORK AND ROLLING FORECAST PROCEDURE

To evaluate the predictive performance of the Bayesian Variable Selection (BVS) model, we implement an out-of-sample forecasting procedure based on a rolling regression framework and a direct prediction approach. This method allows for robust evaluation of forecast accuracy across various horizons and prevents overfitting by dynamically updating the estimation sample.

[3] For example, when $K = 5$ and $J = 2$, the predictor matrix X_t includes the current intercept and 10 lagged predictor values:

$$X_t = \begin{bmatrix} 1 & x_{1,t} & \cdots & x_{5,t} & x_{1,t-1} & \cdots & x_{5,t-1}. \end{bmatrix}$$

The corresponding regression coefficient vector is defined as

$$\beta = \left[\beta_0, \beta_{1,1}, \ldots, \beta_{1,K}, \beta_{2,1}, \ldots, \beta_{2,K}\right]'.$$

The full dataset is divided into two segments: an in-sample period for parameter estimation and an out-of-sample period for forecast evaluation. The last 36 months of the sample are reserved for out-of-sample forecasting (OS_size = 36). At each forecast origin, the model is re-estimated using data available up to that point, and a forecast is generated for future values at multiple horizons.

We conduct forecasting for five different horizons: $h = 1$, 3, 6, 9, and 12, corresponding to one-month, three-month, six-month, nine-month, and twelve-month-ahead predictions, respectively. For each horizon, a separate model is estimated and a direct forecast is produced without relying on recursive steps. This approach enables us to evaluate short-term, medium-term, and long-term forecasting performance independently.

We adopt a 'rolling window strategy' with a fixed in-sample size and sequentially move the forecast origin one period forward. For each iteration and each forecast horizon, the following steps are executed:

(1) The model is estimated via Gibbs sampling using a spike-and-slab prior on each regression coefficient. Posterior samples are drawn for the regression coefficients β_k, inclusion indicators γ_k, prior variances b_0, b_1, model inclusion probability p, and error variance σ^2.

(2) A direct forecast is made for y_{t+h} using the posterior distribution:

$$y_{t+h}^{(s)} = x_{t+h}^{\top}(\beta^{(s)} \cdot \gamma^{(s)}) + \varepsilon^{(s)}, \quad \varepsilon^{(s)} \sim \mathcal{N}(0, \sigma^{2(s)}),$$

where $(\beta^{(s)}, \gamma^{(s)}, \sigma^{2(s)})$ are posterior draws, and x_{t+h} is the predictor vector (including intercept) used for forecasting.

(3) After discarding burn-in samples, the posterior predictive distribution of y_{t+h} is obtained. The point forecast is computed as the posterior mean, and forecast accuracy is evaluated using squared prediction errors and log predictive likelihood.

By repeating this procedure over the 36 out-of-sample periods and across all forecast horizons, we obtain distributions of forecast errors for each h. The Root Mean Squared Error (RMSE) and the log Predictive Probability(log PPL) are calculated to assess the model's predictive accuracy and distributional fit at each forecast horizon. This rolling forecasting framework enables us to rigorously evaluate how the model performs across different temporal horizons and under evolving economic conditions.

3.3. FORECAST EVALUATION METRICS

The accuracy of the forecasts is assessed by comparing them to the actual values observed during the out-of-sample period. Common evaluation metrics include Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE). In this study, we use RMSE as the primary evaluation criterion, as it is widely adopted in macroeconomic forecasting and places greater weight on large errors.

The Root Mean Squared Error for forecast horizon h is computed as:

$$\text{RMSE}(h) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_{t_i+h} - \hat{y}_{t_i+h})^2}.$$

Here, t_i represents the ith forecast origin during the rolling forecast procedure, and y_{t_i+h} denotes the actual value observed h steps ahead.[4] The point forecast $\hat{y}_{t_i+h}$ is calculated as the posterior predictive mean given information available at time t_i.

In addition to RMSE, we also evaluate the quality of the full predictive distribution using the log posterior predictive likelihood (logPPL). This is computed as the sum of log posterior predictive densities over the out-of-sample period:

$$\text{logPPL}(h) = \sum_{i=1}^{N}\log p(y_{t_i+h} \mid \mathcal{F}_{t_i}).$$

A higher value of logPPL indicates better alignment between the predictive distribution and the actual outcomes. While RMSE captures the accuracy of point forecasts, logPPL evaluates the overall calibration and sharpness of the probabilistic forecasts, making it a complementary measure in the Bayesian framework.

3.4. BENCHMARK MODEL

We employ three benchmarks to contextualize the performance of the BVS approach: (i) an autoregressive model AR(p), (ii) a random walk (RW), and (iii) a Bayesian full regression that includes all predictors at all lags without selection. For (iii), to isolate selection from shrinkage, we

[4]For example, if $h = 3$ and the ith forecast origin is t_i = February 2022, then y_{t_i+3} corresponds to the realized value in May 2022. The index i runs from 1 to N, where $N = 36$ is the total number of out-of-sample forecast iterations.

estimate it under the two prior environments used in our BVS specifications, denoted BVS1-Full and BVS2-Full.

To evaluate the predictive performance of the proposed Bayesian Variable Selection (BVS) model, we compare it against a traditional benchmark: the p-th order autoregressive model, AR(p), which is widely used in time series forecasting. The AR(p) model is defined as:

$$y_{t+h} \mid \mathcal{F}_t, \theta \sim \mathcal{N}\left(c + \sum_{j=1}^{p} \phi_j y_{t-j+1},\ \sigma^2\right).$$

Here, y_t denotes the dependent variable at time t, c is the intercept term, and ϕ_j are autoregressive coefficients. The error variance σ^2 is assumed to be constant across forecast horizons.

The lag order p is selected recursively to minimize the Root Mean Squared Error (RMSE) over the out-of-sample forecast window. Specifically, we set a maximum allowable lag P, and for each forecast horizon h, all AR(p) models with $p \leq P$ are estimated. The model with the lowest RMSE is then selected as the optimal benchmark for comparison.

In addition to the autoregressive benchmark, we also consider the random walk (RW) model, a widely used naive benchmark in macroeconomic forecasting. The RW assumes that the best forecast equals the most recently observed value (no-change forecast). Let $y_t = y_{t-1} + u_t$ with $u_t \sim \mathcal{N}(0, \sigma_u^2)$. Then the h-step-ahead predictive distribution is

$$y_{t+h} \mid \mathcal{F}_t \sim \mathcal{N}(y_t,\ h\sigma_u^2),$$

so that the predictive mean is y_t while the variance scales linearly with the horizon h. In implementation, we estimate σ_u^2 from the rolling training window using the first-difference residuals $\{y_t - y_{t-1}\}$, and evaluate density forecasts with $\sigma_h^2 = h\hat{\sigma}_u^2$.[5] This RW benchmark provides a parsimonious yardstick against which the predictive gains of the BVS model can be evaluated, particularly when structural breaks or low signal-to-noise ratios make sophisticated models hard to beat in real time.

To specifically quantify the value of the variable selection mechanism, we additionally consider a Bayesian regression model that includes all available predictors without imposing any selection. Unlike the BVS

[5] Equivalently, one may estimate a horizon-specific error variance σ_h^2 nonparametrically from the training-sample forecast errors $e_{t,h} = y_{t+h} - y_t$. Under homoskedastic innovations both approaches coincide up to sampling noise.

model, where the inclusion indicators $\gamma_{j,k}$ control whether a predictor enters the regression, the full regression benchmark sets all inclusion indicators equal to one, such that every predictor at every lag is forced into the model. The model is therefore specified as:

$$y_{t+h} \mid \mathcal{F}_t, \theta \sim \mathcal{N}\left(\beta_0 + \sum_{j=1}^{J}\sum_{k=1}^{K} X_{k,t-j+1}\beta_{j,k},\ \sigma^2\right),$$

where β_0 is the intercept term; $X_{k,t-j+1}$ is the value of the k-th predictor at lag j; $\beta_{j,k,h}$ are regression coefficients estimated by Bayesian posterior sampling; in this full-regression benchmark all coefficients use *slab-only* Gaussian priors; σ^2 is the variance of the error term.

This benchmark captures the predictive performance of a Bayesian regression without variable selection, thereby reflecting the potential overfitting problem when all predictors are included. Comparing its results against the BVS model highlights the importance of variable selection in improving forecast accuracy and preventing over-parameterization.

4. DATA AND PREDICTOR VARIABLES

This study analyzes the year-on-year growth rate of the China Newly Built Housing Price Index as the dependent variable. The data are monthly and span the period from January 2011 to July 2025. The list of predictor variables is provided in Table 2. A total of 14 predictors are used, grouped into five categories: housing prices, inflation, interest rates, economic fluctuations, and demographics.

To ensure comparability, most predictors are expressed as year-on-year (YoY) growth rates, except for two cases. First, the Real Estate Climate Index (RECI) is transformed to reflect deviations from neutral market conditions (i.e., RECI -100). This form improves interpretability—positive values suggest overheating and negative values imply sluggishness—and yields significantly lower forecast errors.[6] For notational simplicity, we continue to refer to this transformed variable as RECI. Second, the Population Natural Growth Rate (PNGR) is available only at an annual frequency and measured in per mille (‰). To align with the monthly frequency of the forecasting framework, each annual observation

[6] Forecasts using the deviation form of RECI produce lower Root Mean Squared Error (RMSE) values compared to forecasts using its year-on-year growth rate.

N	Data (Growth Rate (%, YoY))	Abbreviation	Category	Source
1	China Newly Built Housing Price Index	HPI	Housing Prices	National Bureau of Statistics of China
2	Real Estate Climate Index	RECI		
3	Consumer Price Index	CPI	Inflation	
4	M2 Money Supply	M2		
5	Global Price of Brent Crude	Oil		Federal Reserve Bank of St. Louis (FRED)
6	Value-added of Industry	VAI	Business Fluctuations	National Bureau of Statistics of China
7	Consumer Confidence Index	CCI		
8	Manufacturing Purchasing Managers Index	PMI		China Federation of Logistics and Purchasing
9	Shanghai SE Composite Index	SSEC		Investing.com
10	10-Year Treasury Bond Yield	Bond10	Interest Rate	China Central Depository and Clearing Co., Ltd.
11	U.S. 10-Year Treasury Bond Yield	USBond10		Investing.com
12	Housing Provident Fund Loan Interest Rate (Loan1-5)	HPF		Beijing Housing Fund Management Center
13	Individual Housing Loan Lending Rate	IHLL		People's Bank of China (CEIC)
14	Population Natural Growth Rate	PNGR	Demographics	National Bureau of Statistics of China

Table 2: List of Predictor Variables. For variables sharing the same data source, the "Source" column is only filled once for brevity. Empty cells in the "Source" column indicate that the source is the same as that of the previous entry.

Figure 1: Housing Price Index (HPI) Year-on-Year Growth Rate. This figure shows the monthly growth rate of China's newly built Housing Price Index (HPI YoY) from January 2011 to July 2025.

is replicated across twelve months. We further standardize PNGR using a z-score transformation, which places the demographic indicator on a comparable scale with the other predictors.[7]

Housing Price Index (HPI) The Housing Price Index (HPI) reflects the overall trend in residential sales prices and serves as a key benchmark for real estate participants. It is compiled monthly by the National Bureau of Statistics (NBS) based on data from 70 major cities, including municipalities and provincial capitals, and captures both first-tier and non-first-tier markets.

As shown in Figure 1, the year-on-year growth rate of HPI remained largely positive between 2011 and mid-2022, with particularly strong growth from 2016 to 2019. Following the COVID-19 outbreak, however, price growth began to slow, turning negative by mid-2022 and declining further through 2023 and 2024. Since December 2024, the trend has slightly recovered, though it is unclear whether this reflects a lasting stabilization.

China's Real Estate Climate Index (RECI) The National Real Estate Climate Index (RECI) is constructed based on the theory of economic cycle fluctuations, drawing upon business cycle theory and analytical methods. It employs time series analysis, multivariate statistics, and

[7]When PNGR is included as a predictor in z-score standardized form, the out-of-sample RMSE of HPI forecasts is noticeably smaller than when using the year-on-year growth rate of PNGR as the predictor.

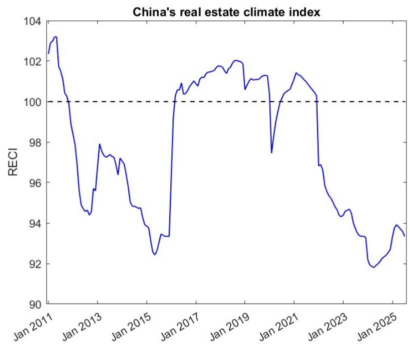

Figure 2: China's Real Estate Climate Index (RECI). This figure shows the monthly Real Estate Climate Index (RECI) for China from January 2011 to July 2025. The index is normalized to 100 in 2012, with values between 95 and 105 indicating a moderate market climate, below 95 a depressed market, and above 105 an overheated market.

econometric techniques, with real estate development investment serving as the benchmark indicator. The index incorporates a set of variables related to real estate investment, funding, construction area, and sales, with seasonal effects removed and stochastic components included. It is compiled using a growth rate cycle method and is revised monthly as new data become available. The year 2012 is designated as the base year, with the index level set to 100. In general, a RECI value of 100 represents the optimal climate level; values between 95 and 105 indicate a moderate climate; values below 95 suggest a relatively depressed climate; and values above 105 reflect an overheated market condition. Overall, the RECI provides an early warning signal for macroeconomic regulation and serves as a valuable statistical tool to support investment decision-making. The historical data from January 2011 to July 2025 are shown in Figure 2.

From Figure 2, we can observe that during the period of particularly strong housing price growth between 2016 and 2019, the RECI index also remained at a relatively high level. When HPI YoY turned negative in early 2022, the RECI likewise experienced a rapid decline. After December 2024, the index value remained below 95 but showed a mild rebound. These turning points align closely with the critical shifts in the HPI YoY, suggesting that RECI may serve as an important predictor of housing price dynamics.

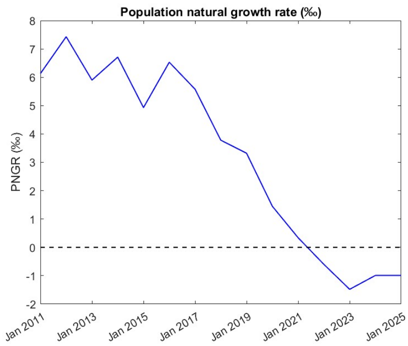

Figure 3: POPULATION NATURAL GROWTH RATE (PNGR). This figure shows the annual population natural growth rate of China over the period January 2011 to July 2025.

Housing Provident Fund Loan Interest Rate (HPF) Housing Provident Fund Loan Interest Rate (HPF) refers to the interest rate applied to homebuyers who obtain loans through China's Housing Provident Fund system. This rate is centrally determined and adjusted by the People's Bank of China (PBoC), based on macroeconomic conditions such as economic growth and monetary policy objectives. As a form of national policy-based lending, Housing Provident Fund loans are available exclusively to employees who contribute to the fund. These loans offer lower interest rates compared to commercial mortgages and exhibit relatively greater stability over time. Interest rates may vary depending on loan maturity and repayment terms. In practice, the loan interest rate is divided into two categories: short-term (1–5 years) and long-term (6–30 years).[8]

Individual Housing Loan Lending Rate (IHLL) Individual Housing Loan Lending Rate (IHLL) refers to the interest rate charged on commercial personal housing loans in China. It is a major component of the national interest rate system and directly affects homebuyers' borrowing costs. In recent years, China has promoted interest rate marketization by shifting the pricing benchmark for personal housing loans from the traditional benchmark rate to the Loan Prime Rate (LPR), which is adjusted

[8]In this study, we use only the short-term 1–5 year Housing Fund Loan interest rate as the predictor variable. Description of this policy-based loan are based on information from Baidu Baike.

monthly.

Population natural growth rate (PNGR) The population natural growth rate, abbreviated as PNGR, refers to the ratio of the natural increase in population (births minus deaths) to the average total population over a given period (typically one year), expressed in per mille (‰).

Figure 3 illustrates the historical evolution of PNGR from 2011 to 2025. The rate exhibits a persistent downward trend, declining from above 6‰ in the early 2010s to negative values after 2022, reflecting China's structural demographic transition. This demographic shift coincides with weakening housing demand fundamentals, suggesting that PNGR has substantial implications for the long-term forecasting of housing prices.

5. EMPIRICAL ANALYSIS

5.1. PRIORS

As discussed in Section 2, the choice of priors is critical in Bayesian variable selection, as the hyperparameters α_{00}, δ_{00}, α_{01}, and δ_{01} determine the prior variances b_0 (spike) and b_1 (slab). These govern the shrinkage pattern: a smaller b_0 enforces stronger shrinkage on coefficients associated with irrelevant predictors, while a larger b_1 grants more flexibility for relevant coefficients. Poorly specified values may result in excessive shrinkage or insufficient sparsity. To assess the impact of shrinkage structure on forecast performance, we consider two prior settings—BVS1 and BVS2—that differ in their spike–slab variance scales.

The priors for BVS1 are specified as:

$$a_0 = 5, \quad c_0 = 5,$$

$$\alpha_{00} = \alpha_{01} = 10, \quad \delta_{00} = 0.03^2 \times 10, \quad \delta_{01} = 10,$$

$$\alpha_0 = 20, \quad \delta_0 = 4.$$

The priors for BVS2 are specified as:

$$a_0 = 5, \quad c_0 = 5,$$

$$\alpha_{00} = \alpha_{01} = 10, \quad \delta_{00} = 0.01^2 \times 10, \quad \delta_{01} = 100,$$

$$\alpha_0 = 20, \quad \delta_0 = 4.$$

Under these specifications, BVS2 induces a tighter spike (b_0 smaller) and a looser slab (b_1 larger), promoting sparser models by shrinking irrelevant coefficients more aggressively while allowing selected coefficients greater variability. In contrast, BVS1 features a relatively looser spike and a tighter slab, yielding denser models with more conservative magnitudes for selected coefficients.

We evaluate out-of-sample forecasting performance for both BVS1 and BVS2 across multiple horizons using two complementary criteria: (i) the root mean squared error (RMSE) for point forecasts, and (ii) the log predictive probability (log PPL) for density forecasts. Comparing results under both metrics allows us to identify which prior configuration is better suited to forecasting China's housing price dynamics at each horizon.

5.2. BAYESIAN VARIABLE SELECTION MODEL: PREDICTIVE DISTRIBUTION AND FORECAST EVALUATION

We evaluate the forecast performance of the Bayesian variable selection (BVS) model using monthly data spanning January 2011 to July 2025, yielding a total sample of 174 observations. Forecasts are generated with a rolling regression scheme under a direct multi-horizon setting. The out-of-sample forecasting window is fixed at 36 months, covering all periods after the structural change in the year-on-year growth rate of the housing price index. As benchmarks, we consider the random walk (RW), autoregressive models of order p [AR(p)], and a full-regression BVS specification that forces all predictors into the model (BVS-Full). For comparability across specifications, all models are estimated using an MCMC sampler with 11,000 iterations. We discard the first 1,000 as burn-in and retain 10,000 posterior draws.

Forecast performance is evaluated along two complementary dimensions: (i) point forecasts measured by root mean squared error (RMSE), and (ii) density forecasts assessed by log posterior predictive likelihood (Log PPL). The evaluation is conducted across forecast horizons $h \in \{1, 3, 6, 9, 12, 18, 24\}$. For the AR($p$) benchmark, we allow up to 24 monthly lags ($p \leq 24$) to capture annual and bi-annual dynamics and the slow adjustment observed in housing markets.[9] By contrast, in the multivariate BVS specification, the lag length is restricted to at most 3 for pragmatic

[9]See Case and Shiller (1989) and Genesove and Mayer (2001) for housing market inertia.

RMSE Lag = 1

h	**RW**	AR(p)	BVS1-Full	BVS1	BVS2-Full	BVS2
1	0.40	**0.16**	0.39	1.22	0.39	0.91
3	1.14	0.68	0.93	2.04	0.87	1.18
6	2.13	1.75	1.18	1.95	1.12	1.25
9	2.90	2.92	1.98	1.85	2.03	1.83
12	3.43	4.12	3.02	**2.41**	2.73	2.52
18	4.05	5.81	2.65	2.25	2.21	2.34
24	4.35	6.17	2.94	**2.58**	2.92	2.94

RMSE Lag = 2

h	**RW**	AR(p)	BVS1-Full	BVS1	BVS2-Full	BVS2
1	0.40	**0.16**	0.17	0.18	0.19	0.17
3	1.14	0.68	0.57	0.92	**0.54**	0.55
6	2.13	1.75	1.10	0.94	1.14	0.89
9	2.90	2.92	2.25	**1.70**	2.11	1.76
12	3.43	4.12	3.32	2.94	2.89	2.70
18	4.05	5.81	2.84	1.75	2.37	1.90
24	4.35	6.17	3.93	3.28	3.81	4.24

RMSE Lag = 3

h	**RW**	AR(p)	BVS1-Full	BVS1	BVS2-Full	BVS2
1	0.40	**0.16**	0.17	0.18	0.18	0.17
3	1.14	0.68	0.57	1.02	**0.54**	0.56
6	2.13	1.75	1.26	**0.88**	1.22	0.96
9	2.90	2.92	2.46	1.91	2.19	1.82
12	3.43	4.12	3.42	2.96	2.92	2.71
18	4.05	5.81	2.98	**1.67**	2.48	**1.67**
24	4.35	6.17	4.69	4.00	4.19	4.67

Table 3: Out-of-sample RMSE comparison across models by lag. RW = random walk; "Full" = all predictors included. AR(p): minimum RMSE over $p \leq 24$ at each horizon. Bold numbers indicate, at each horizon, the lowest RMSE across all models and lag specifications.

Log PPL Lag = 1						
h	**RW**	AR(p)	BVS1-Full	BVS1	BVS2-Full	BVS2
1	-26.18	**4.43**	-21.67	-77.13	-21.96	-56.51
3	-56.17	-40.39	-54.51	-77.16	-53.47	-59.01
6	-80.53	-71.50	-67.05	-75.49	-66.83	-68.18
9	-93.89	-89.79	-77.41	-76.91	-77.84	-76.11
12	-100.35	-102.61	-90.33	**-82.61**	-87.54	-84.93
18	-104.37	-117.32	-85.22	-80.70	-79.97	-81.91
24	-104.98	-121.05	-86.50	**-84.37**	-87.98	-88.64
Log PPL Lag = 2						
h	**RW**	AR(p)	BVS1-Full	BVS1	BVS2-Full	BVS2
1	-26.18	**4.43**	0.69	-0.35	-0.07	1.26
3	-56.17	-40.39	-39.41	-49.43	**-38.89**	-39.93
6	-80.53	-71.50	-57.73	-57.64	-58.28	**-55.44**
9	-93.89	-89.79	-79.47	**-71.23**	-76.92	-71.67
12	-100.35	-102.61	-96.56	-90.16	-89.73	-88.18
18	-104.37	-117.32	-87.25	-73.90	-81.13	-76.21
24	-104.98	-121.05	-100.54	-91.63	-103.30	-111.53
Log PPL Lag = 3						
h	**RW**	AR(p)	BVS1-Full	BVS1	BVS2-Full	BVS2
1	-26.18	**4.43**	-0.27	-0.35	-0.00	1.43
3	-56.17	-40.39	-39.75	-52.38	-39.01	-40.19
6	-80.53	-71.50	-59.33	-56.23	-59.13	-55.73
9	-93.89	-89.79	-82.53	-73.75	-77.86	-72.67
12	-100.35	-102.61	-97.68	-89.94	-89.21	-88.09
18	-104.37	-117.32	-90.21	**-72.62**	-82.78	-72.68
24	-104.98	-121.05	-114.06	-101.85	-107.51	-117.49

Table 4: Out-of-sample Log PPL comparison across models by lag. RW = random walk. "Full" denotes the BVS full-regression variant (all predictors forced in). AR(p): maximum Log PPL over $p \leq 24$ at each horizon. Bold numbers indicate the highest Log PPL across all models and lag specifications at each forecast horizon.

reasons: exploratory out-of-sample checks with lags 4–5 yield no RMSE gains, and a shorter lag structure curbs model-space proliferation and to enable clear visualization and interpretation of key predictors in Section 5.3.

As shown in Tables 3 and 4, the BVS models (BVS1 and BVS2) outperform the random walk and AR(p) benchmarks at all forecast horizons except $h = 1$; for $h \geq 3$, they deliver lower RMSEs and higher Log PPL. At $h = 1, 3$, BVS-Full performs comparably to the sparse BVS models; for $h \geq 6$, the sparse BVS tends to outperform, reflecting parsimony in the predictor set. Between priors, BVS2 (tighter spike, looser slab) tends to lead in short-horizon performance, whereas BVS1 delivers more stable forecast accuracy at medium to long horizons. A multi-lag design is beneficial; at short-to-medium horizons, models with two or three lags generally outperform single-lag specifications.

Guided by Tables 3 and 4, we follow a simplicity-first rule: at each horizon, we choose the simplest model that performs competitively on both RMSE and Log PPL. To enhance visualization and interpretability, we restrict each predictor to at most two lags when identifying key predictors. At $h = 1$, the differences in RMSE and Log PPL between BVS2 with lag 2 and lag 3 are negligible. Similarly, at $h = 18$, BVS1 with lag 2 and lag 3 shows only minor differences, further supporting the two-lag restriction. To ensure forecast accuracy, however, we relax the lag cap (lags 1–3) and select the horizon-specific optimum based on both metrics. Accordingly, we report two selection standards—one for influencing-factor analysis (lags up to two) and one for forecasting accuracy (allowing up to three)—as summarized in Table 5.

Horizon (h)	**Influencing-factor analysis**		**Forecasting accuracy**	
	Prior	Lag	Prior	Lag
1	BVS2	2	BVS2	2
3	BVS2	2	BVS2	2
6	BVS2	2	BVS2	2
9	BVS1	2	BVS1	2
12	BVS1	1	BVS1	1
18	BVS1	2	BVS1	3
24	BVS1	1	BVS1	1

Table 5: Model selection by purpose and forecast horizon. BVS1 and BVS2 are the prior settings in Section 5.1. Selections follow the parsimony-first rule: up to two lags for the influencing-factor analysis; up to three for forecasting and comparing out-of-sample predictive distributions.

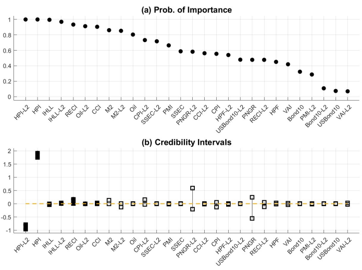

(a) $h = 1$

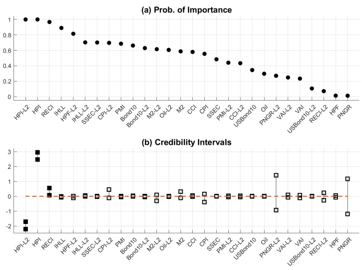

(b) $h = 3$

Figure 4: Posterior inclusion probabilities and credibility intervals (short-term forecasts). Variables labeled with “L2” denote lag-2 predictors; unlabeled variables (e.g., HPI) refer to lag-1 predictors.

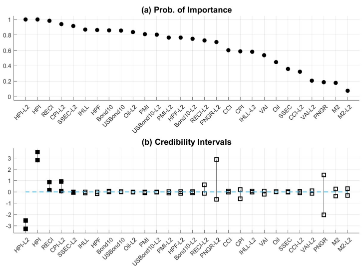

(a) $h = 6$

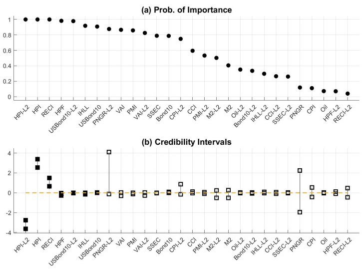

(b) $h = 9$

Figure 5: Posterior inclusion probabilities and credibility intervals (mid-term forecasts). Variables labeled with “L2” denote lag-2 predictors; unlabeled variables (e.g., HPI) refer to lag-1 predictors.

5.3. KEY PREDICTORS OF CHINA’S HOUSING PRICE GROWTH

This section investigates the key macroeconomic and financial variables associated with the year-on-year growth rate of China’s newly built housing price index (HPI), using results from the Bayesian Variable Selection (BVS) model. The analysis is conducted separately for each forecast horizon ($h = 1, 3, 6, 9, 12, 18, 24$), following the optimal model and lag specification identified in Section 5.2. For posterior inference, we run a Markov chain Monte Carlo (MCMC) sampler for a total of 11,000 iterations, discarding the first 1,000 as burn-in.

Figures 4 to 6 summarize the findings. Each figure consists of two panels: Panel (a) displays predictors ranked by posterior inclusion proba-

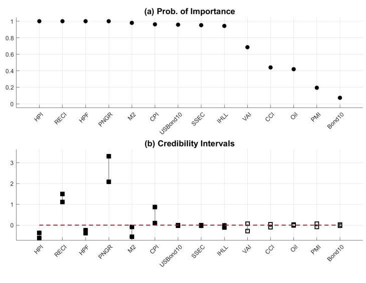

(a) $h = 12$

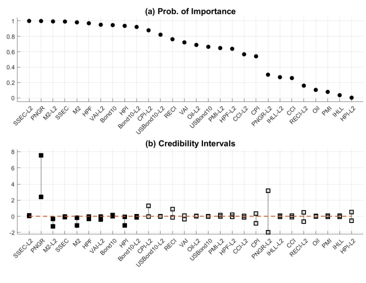

(b) $h = 18$

Figure 6: POSTERIOR INCLUSION PROBABILITIES AND CREDIBILITY INTERVALS (LONG-TERM FORECASTS). Variables labeled with "L2" denote lag-2 predictors; unlabeled variables (e.g., HPI) refer to lag-1 predictors.

bilities, and Panel (b) shows the corresponding 90% posterior credible intervals for the regression coefficients. The posterior probability of $\gamma_k = 1$ corresponds to the mean of the posterior samples for γ_k, reflecting the likelihood of a predictor being included. Black boxes represent credible intervals that exclude zero, indicating statistical significance at the 90% level, while white boxes indicate non-significance. To highlight how predictor relevance varies by forecast horizon, we group the analysis into four buckets: short term ($h = 1,3$), medium term ($h = 6,9$), long term ($h = 12,18$), and ultra-long term ($h = 24$).

Short-term key predictors ($h = 1,3$) Short-term housing price dynamics are primarily driven by strong momentum effects. Both the current HPI and its second lag (HPI-L2) consistently exhibit the highest posterior inclusion probabilities and statistically significant coefficients, confirming that recent price trends are key determinants of near-term fluctuations. In addition, the Real Estate Climate Index (RECI) also appears as an influential factor, suggesting that real estate market sentiment and cyclical conditions contribute meaningfully to short-term price movements. These results indicate that short-term housing price growth in China is mainly governed by internal price inertia and prevailing market conditions, as illustrated in Figure 4.

Medium-term key predictors ($h = 6,9$) Over the medium term, the importance of RECI continues to rise, indicating that promoting real-

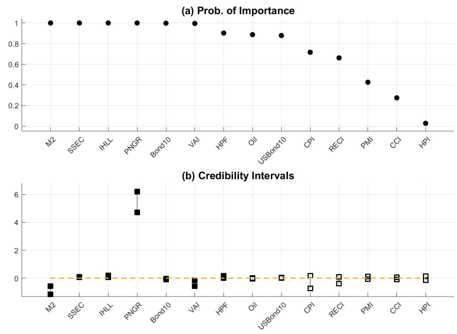

Figure 7: Posterior inclusion probabilities and credibility intervals (ultra-long forecasts). Variables labeled with "L2" denote lag-2 predictors; unlabeled variables (e.g., HPI) refer to lag-1 predictors.

estate investment plays an important role in medium-term increases in housing prices ($h = 6$–9). Inflation (CPI) emerges as significant with a positive association, consistent with inflation-hedging demand for housing. Lagged HPI growth remains a key predictor, capturing momentum effects. Policy-related credit indicators also matter: the Housing Provident Fund (HPF) lending rate and the IHLL index enter with negative coefficients, implying that tighter mortgage and credit conditions may suppress price growth, whereas lower borrowing costs would ease downward pressures. Taken together, medium-term forecasts are jointly shaped by macroeconomic fundamentals (inflation and market activity) and the credit-policy stance, as illustrated in Figure 5.

Long-term key predictors ($h = 12, 18$) At long horizons, demographics become central. The population natural growth rate (PNGR) carries a high positive median effect; at $h = 12$ it already ranks among the top predictors, and by $h = 18$ it becomes the single most influential factor for housing price growth. At $h = 12$, RECI and CPI remain important with positive contributions, while the Housing Provident Fund (HPF) lending rate exhibits a stronger negative effect than in the medium term, indicating that tighter mortgage credit continues to restrain price growth at longer horizons. At $h = 18$, HPF remains relevant and broad money (M2) enters with a small negative coefficient. The negative sign on M2 can be interpreted as a countercyclical policy-response effect, acting as a "bad-times" signal rather than a pure liquidity driver, as illustrated in Figure 6.

Horizon (h)	Influencing Factors
1	HPI, RECI
3	HPI, RECI
6	CPI, RECI, HPI
9	RECI, HPI, –HPF, –IHLL
12	PNGR, RECI, CPI, –HPF, –M2, –IHLL
18	PNGR, –M2, –HPF
24	PNGR, –M2

Table 6: SUMMARY OF KEY PREDICTORS ACROSS FORECAST HORIZONS. This table summarizes the main predictors selected by the Bayesian Variable Selection (BVS) model for each forecast horizon h. The ordering of variables reflects their relative importance. By convention, positive associations are not marked, whereas a minus sign indicates that the predictor is negatively related to housing price growth.

Ultra-long term key predictors ($h = 24$) At the ultra-long horizon, the population natural growth rate (PNGR) emerges as the nearly exclusive driver of housing price growth. In addition, M2 also appears with a small negative coefficient, which can be interpreted as reflecting its role as a countercyclical policy-response indicator. The effects of other variables become negligible at this horizon, as illustrated in Figure 7.

To provide a concise overview of how predictor importance evolves across forecast horizons, Table 6 summarizes the key variables identified by the BVS model for each value of h. The ordering of variables indicates their relative importance, and the positive/negative signs denote the direction of effects whose posterior credible intervals exclude zero.

5.4. OUT-OF-SAMPLE FORECAST DISTRIBUTIONS

For an intuitive visualization of model performance, Figures 8 and 9 display the out-of-sample forecast distributions of the best-performing Bayesian Variable Selection (BVS) models alongside two benchmark models: the random walk (RW) model and the best-performing AR(p) model. Forecast horizons span from short- to medium-term ($h = 1, 3, 6, 9$) and long- to ultra-long-term ($h = 12, 18, 24$). Each panel plots the forecast median (black line), 80% and 95% credible intervals (shaded areas), and the actual HPI year-on-year growth (blue line), enabling transparent compar-

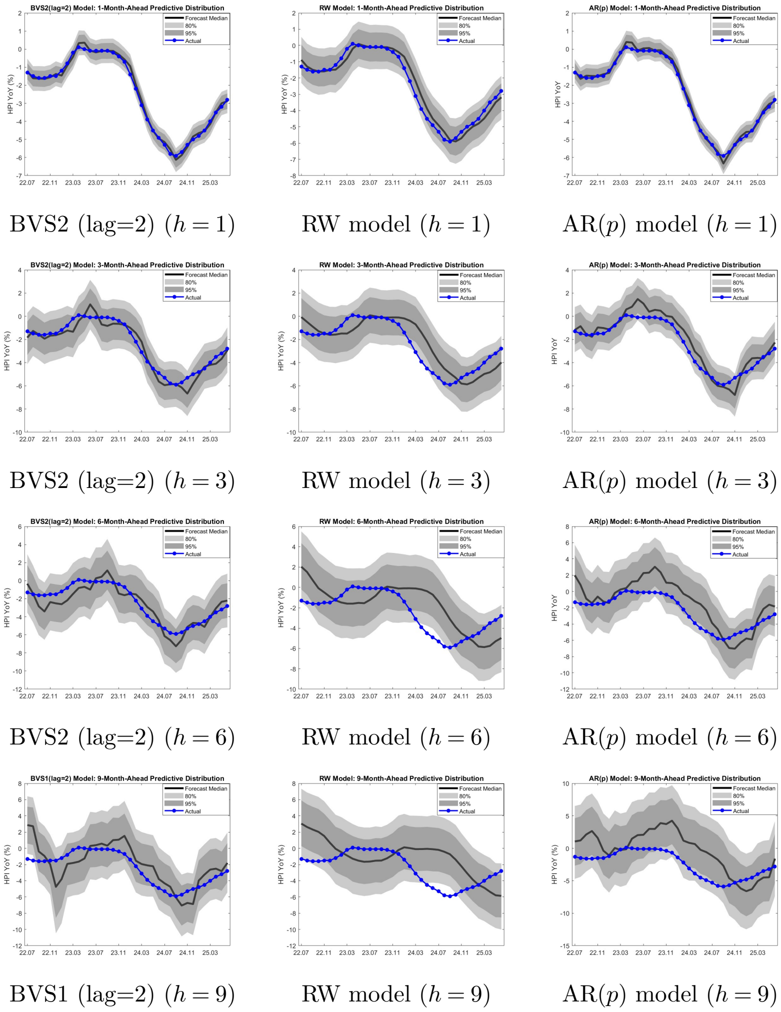

Figure 8: Out-of-sample forecast distributions of HPI YoY for $h = 1, 3, 6, 9$ across three models (left to right). Best-performing BVS, RW benchmark, and best-performing AR(p).

isons of predictive accuracy and uncertainty across models and horizons.

Across all horizons, the BVS models demonstrate superior predictive performance compared to both the RW and AR(p) benchmarks. The forecast medians from the BVS models more closely track the actual housing

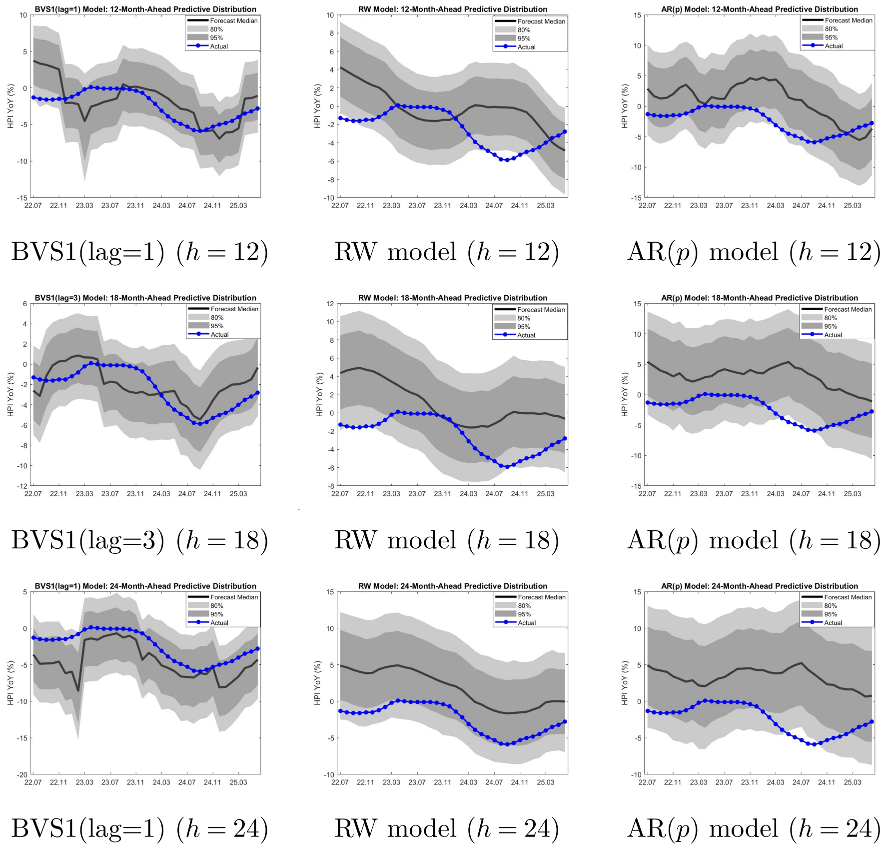

Figure 9: Out-of-Sample Forecast Distributions of HPI YoY for $h = 12, 18, 24$ across tree models (left to right). Best-performing BVS, RW benchmark, and best-performing AR(p).

price growth, and the associated credible intervals are generally narrower, indicating improved uncertainty quantification. One exception is at the 1-month horizon ($h = 1$), where the AR(p) model exhibits slightly better forecast accuracy than BVS. These results are consistent with the RMSE and log PPL comparisons reported in Section 5.2.

Forecast performance varies by horizon. At short horizons ($h = 1, 3$), BVS models closely match the actual trajectory with narrow uncertainty bands. As the horizon increases ($h = 6, 9$), the forecast medians remain close to the actual values and lie within the 95% credible intervals. At $h = 12, 18, 24$, forecast precision diminishes as the credible intervals widen

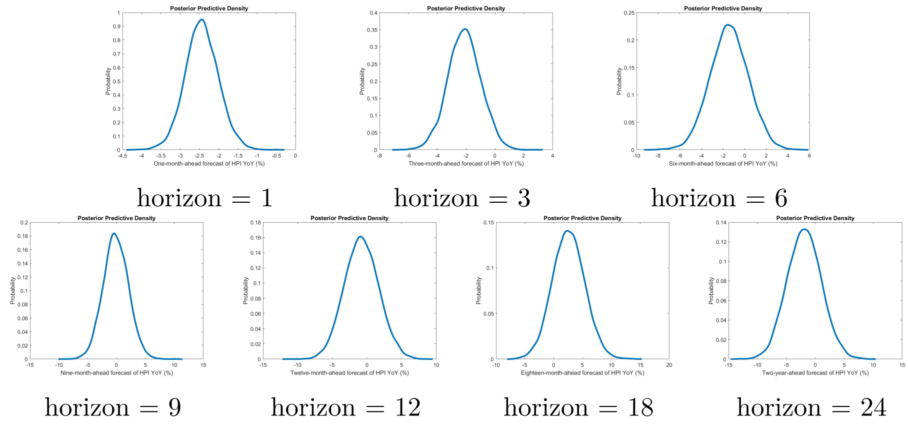

Figure 10: Posterior forecast densities of HPI YoY at horizons $h = 1$–24. This figure shows the forecast distributions of China's housing price growth (HPI YoY) from one month to 24 months ahead. Each distribution is obtained from the best-performing Bayesian Variable Selection (BVS) model at the corresponding horizon, as identified by the forecasting accuracy criteria in Table 4.

further, and actual values occasionally fall outside the 95% bands. In the most recent periods of the out-of-sample window, the forecast medians show a slight upward bias at horizons $h = 12$ and $h = 18$, while at $h = 24$ the median is marginally lower than the actual values. This pattern suggests a tendency to overestimate housing price growth at 12–18 months and underestimate it at 24 months. Importantly, actual values remain within the 95% credible intervals, underscoring the robustness of the BVS model even at longer horizons.

5.5. FORECAST DISTRIBUTIONS FOR HPI YOY

In this section, we generate 24-month-ahead forecasts of the housing price index (HPI YoY) using the direct forecasting approach. For each horizon, the forecasting model is selected according to the horizon-specific optimal specifications presented in Table 5 of Section 5.2.

The Bayesian framework enables full probabilistic forecasting, allowing us to obtain not only point predictions but also entire forecast distributions. This provides a richer description of future uncertainty and improves the interpretation of risks associated with housing market dy-

Horizon (h)	Quantile 5%	Median 50%	Quantile 95%
1	-3.15	-2.45	-1.72
3	-4.00	-2.12	-0.22
6	-4.32	-1.44	1.50
9	-3.67	-0.18	3.46
12	-4.75	-0.80	3.29
18	-1.93	2.54	7.24
24	-6.71	-1.92	2.94

Table 7: POSTERIOR QUANTILES OF THE HPI YOY FORECASTS. This table reports the posterior predictive quantiles of China's housing price growth (HPI YoY) for forecast horizons from 1 to 24 months.

namics. Figure 10 presents the posterior predictive densities of HPI YoY across horizons from $h = 1$ to $h = 24$, capturing both the central tendency and the uncertainty bands. Forecasts are made for the period from August 2025 to July 2027.

Table 7 reports the 5%, median, and 95% posterior quantiles of the HPI YoY forecasts for each horizon.

Figure 11 presents the posterior median and 90% credible interval of HPI YoY forecasts from August 2025 to July 2027. The median forecast shows a mild upward trend in the first year; however, the pace of recovery is limited, and YoY growth is likely to remain negative for most of the forecast horizon. The forecast shows greater variability between the 12- and 24-month horizons. As discussed in Section 5.4, the forecasts at $h = 12$ and $h = 18$ may be slightly upward biased, while the $h = 24$ forecast may be slightly downward biased. Accounting for these biases, the projected recovery appears remain weaker, with a high probability of continued negative HPI growth through mid-2027.

These results suggest a subdued outlook for China's housing market over the next two years. Investors should remain cautious, and policymakers may consider proactive measures informed by the key macroeconomic drivers identified in Section 5.3.

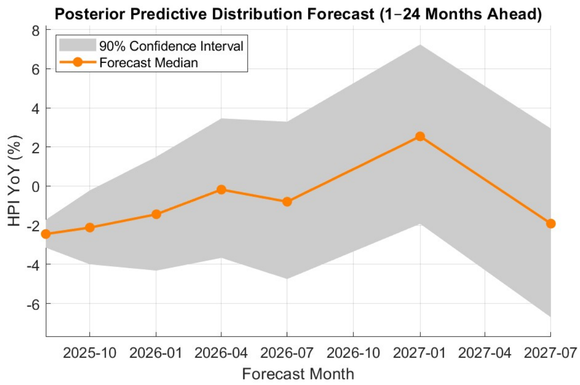

Figure 11: Forecasts of HPI YoY from august 2025 to july 2027. This figure shows the posterior predictive forecasts of China's housing price growth (HPI YoY) for horizons from one to 24 months ahead.

6. BAYESIAN VARIABLE SELECTION COMPARISON WITH RANDOM FOREST

In addition to traditional benchmarks such as AR(p) models and the Random Walk, we also compare the Bayesian Variable Selection (BVS) framework with a modern machine learning method: the Random Forest (RF). This provides a more demanding nonlinear benchmark and helps assess whether the predictive gains of BVS are robust beyond linear time-series models.

6.1. RANDOM FOREST MODEL STRUCTURE

Random Forest, introduced by Breiman (2001), is an ensemble learning method that combines predictions from multiple regression trees. Each tree is trained on a bootstrap sample of the data, and at each split only a random subset of predictors is considered. The forecast is obtained by averaging across all trees, which reduces variance and mitigates overfitting.

Formally, the h-step-ahead Random Forest forecast is given by:

$$\hat{y}_{t+h}^{RF} = \frac{1}{B}\sum_{b=1}^{B} T_b(X_t),$$

where $T_b(\cdot)$ denotes the prediction of the b-th regression tree trained on a resampled dataset, B is the number of trees in the forest, and X_t is the predictor vector at time t (including lagged variables).

Unlike the BVS model, which performs Bayesian variable selection through $\gamma_{j,k}$ indicators, RF implicitly selects variables at each split and models non-linear interactions between predictors. This allows RF to flexibly capture complex dynamics, but it does not provide posterior distributions or probabilistic forecasts.

6.2. FORECASTING PROCEDURE

To ensure comparability with BVS, the Random Forest model is implemented under the same out-of-sample forecasting framework:

(1) We adopt a *direct forecasting* approach for each horizon $h = 1, 3, 6, 9, 12, 18, 24, 30, 36$, where a separate model is trained to predict y_{t+h} directly.

(2) Five lag structures are considered ($J = 1, 2, 3, 4, 5$). For example, under $J = 5$ it expands to $[1, X_{t-4}, X_{t-3}, X_{t-2}, X_{t-1}, X_t]$.

(3) A rolling forecasting window with out-of-sample size of 36 months (August 2022–July 2025) is used. At each forecast origin, the RF model is re-trained on all available in-sample data, and a prediction is generated for y_{t+h}.

(4) Forecast accuracy is evaluated using the Root Mean Squared Error (RMSE), identical to the metrics applied in the BVS, AR(p), and Random Walk benchmarks.

The Random Forest implementation uses 500 trees, a minimum leaf size of 5, and $\sqrt{K}$ predictors sampled at each split, where K denotes the total number of candidate predictors (all variables and their lags). These hyperparameters follow common practice in macro-financial forecasting and ensure stable performance.

6.3. BVS COMPARISON WITH RANDOM FOREST

The two models are evaluated by out-of-sample RMSE across horizons and lag specifications, as reported in Table 8. The results show that BVS dominates across short-, medium-, and long-term horizons. Specifically, for $h = 1, 3, 6, 9, 12, 18$, the Best BVS specification achieves the lowest RMSE in all cases, often with a substantial margin over RF. This indicates

h	**Best BVS**	**RF (lag1)**	**RF (lag2)**	**RF (lag3)**	**RF (lag4)**	**RF (lag5)**
1	**0.17**	1.35	1.07	1.08	1.10	1.16
3	**0.55**	1.71	1.51	1.52	1.54	1.58
6	**0.88**	2.45	2.20	2.19	2.14	2.17
9	**1.70**	2.79	2.59	2.52	2.51	2.56
12	**2.41**	3.13	2.97	2.93	2.91	2.86
18	**1.67**	3.66	3.33	3.20	3.00	2.95
24	2.58	3.20	2.84	2.67	2.53	**2.47**
30	3.64	3.21	2.83	2.71	2.63	**2.48**
36	2.65	3.12	2.67	**2.57**	2.62	2.60

Table 8: Out-of-sample RMSE comparison: BVS vs. Random Forest. The evaluation period spans from August 2022 to July 2025. Each column reports RMSE values of the Best BVS model and Random Forest (RF) under different lag specifications. Bolded values indicate the best performance (lowest RMSE) for each forecast horizon. "Best BVS" denotes the BVS1 or BVS2 model with lag ≤ 3 that achieves the lowest RMSE at each horizon.

that the linear BVS framework with parsimonious predictor selection is highly effective in capturing housing price dynamics within two years. In the ultra-long horizons, beyond two years, the relative performance shifts. For $h = 24$ and $h = 30$, RF with higher lag orders (lag 4–5) outperforms BVS, while at $h = 36$ the best RF specification (lag 3) slightly improves upon BVS. These results suggest that the nonlinear structure of RF may capture longer-term dependencies that BVS does not fully exploit.

The comparative findings are consistent with the broader literature on the relative performance of Bayesian shrinkage methods and machine learning models. Studies such as Koop and Korobilis Koop and Korobilis (2012) and Groen, Paap, and Ravazzolo Groen, Paap and Ravazzolo (2013) show that Bayesian variable selection and shrinkage approaches are particularly effective at short horizons, where efficient elimination of irrelevant predictors enhances forecast precision. In contrast, flexible machine learning models such as random forests and boosting often gain strength at longer horizons by accommodating complex lag structures and nonlinear interactions Medeiros et al. (2019). Overall, while RF provides some value at very long horizons, the BVS framework proves more effective across most relevant horizons in this study.

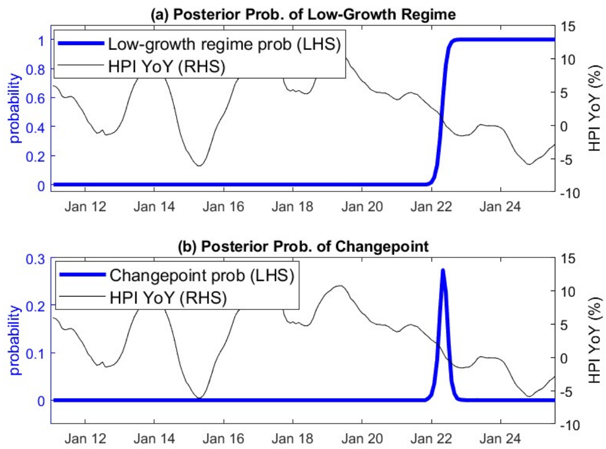

Figure 12: Change point analysis of HPI YoY. This figure shows the posterior probability of the low-growth regime (top panel) and the posterior probability of a change point (bottom panel) in China's housing price growth (HPI YoY), indicating a structural break around March–April 2022.

7. ROBUSTNESS ANALYSIS: STRUCTURAL BREAK IN NATIONAL HOUSING PRICES

To assess the robustness of our forecasting results, we examine whether China's housing price dynamics experienced structural shifts prior to the out-of-sample evaluation. Detecting such breaks is essential for interpreting recent model performance within the broader macroeconomic context. For this purpose, we apply a Bayesian change point model to the HPI YoY series to identify potential regime changes and evaluate whether the recent downturn reflects a persistent structural shift.

We assume the HPI YoY series $\{y_t\}$ follows a regime-switching normal distribution:

$$y_t \mid c_{S_t}, \sigma^2_{S_t} \sim \text{Normal}(c_{S_t}, \sigma^2_{S_t}),$$

where $S_t \in \{1, 2\}$ denotes the latent state at time t, governed by a two-state Markov chain. The prior distributions are:

$$p_{11} \sim \text{Beta}(100, 4), \quad p_{22} = 1,$$

$$c_1, c_2 \sim \text{Normal}(3, 25), \quad \sigma^2_1, \sigma^2_2 \sim \text{InverseGamma}(5, 35).$$

Variables	Mean	s.e.	2.5%	Median	97.5%	Ineff	p-val
$c_{S_t=1}$	4.56	0.38	3.82	4.56	5.31	0.76	0.963
$c_{S_t=2}$	-2.21	0.38	-2.96	-2.21	-1.47	1.50	0.675
$\sigma^2_{S_t=1}$	19.19	2.30	15.10	19.02	24.17	0.78	0.795
$\sigma^2_{S_t=2}$	5.08	1.07	3.34	4.95	7.49	1.34	0.849
p	0.979	0.009	0.958	0.980	0.993	0.91	0.914

Table 9: Bayesian estimation results for HPI YoY. This table reports posterior estimates from a two-state Bayesian regime-switching model for China's national housing price index (HPI YoY). The two regimes reflect a high-growth state ($S_t = 1$) and a low-growth state ($S_t = 2$), with distinct mean growth rates c_{S_t}. The variance parameters $\sigma^2_{S_t}$ indicate greater volatility in the high-growth regime. The transition probability p corresponds to p_{11}, the probability of remaining in regime 1.

Figure 12 illustrates the posterior probability of the low-growth regime and the inferred change point in the HPI YoY series. The results suggest a clear structural break around March–April 2022, when the probability of a regime switch rises sharply and persists throughout the subsequent period.

The Bayesian estimation results (Table 9) implies that the HPI YoY series shifted from a high-growth regime ($\sim 4.6\%$) to a low-growth regime ($\sim -2.2\%$) around March–
April 2022. This structural break coincides with major developments in China's housing policy and macroeconomic environment. In particular, the "three red lines" policy, introduced in 2020, placed strict leverage constraints on real estate developers. Although implemented earlier, its effects accumulated over time, limiting refinancing capacity and worsening the liquidity strains of highly indebted firms such as Evergrande. At the same time, the COVID-19 pandemic depressed household income growth and housing demand, while demographic shifts—especially the declining natural population growth rate—further weighed on the housing market. Together, these factors contributed to the regime transition, marked by weaker price growth and heightened volatility.

The out-of-sample forecasting period in this study (August 2022 to July 2025) largely falls within this low-growth regime. The superior performance of the BVS model—particularly in quantifying forecast uncertainty—demonstrates its usefulness in capturing housing market dynamics under conditions of structural adjustment and financial stress. This robustness check supports the validity of our rolling forecast framework and provides a consistent context for interpreting the empirical results. The next section turns to the broader policy implications of these findings.

8. CONCLUSION AND POLICY IMPLICATIONS

This paper applies a Bayesian Variable Selection (BVS) framework to forecast the year-on-year growth rate of China's newly built housing price index (HPI YoY). Using rolling-window estimation and direct multi-horizon forecasts, we compare BVS with traditional benchmarks and a machine learning alternative, the Random Forest (RF). The results show that BVS consistently outperforms autoregressive (AR) models and the random walk (RW) across most horizons in both point and density forecasts. Against RF, BVS dominates within two years ($h = 1$ to $h = 18$), while RF performs slightly better at very long horizons ($h = 24, 30, 36$). Overall, BVS emerges as the more reliable forecasting tool in our study.

The horizon-specific determinants of housing prices provide valuable economic interpretations. In the short run ($h = 1, 3$), housing price momentum (lagged HPI) and real estate market sentiment (RECI) dominate. In the medium term ($h = 6$ to $h = 9$), inflation (CPI), RECI, and credit conditions (HPF, IHLL) become influential, reflecting the importance of macro-financial interactions. Over the long run ($h = 12, 18, 24$), demographic fundamentals (PNGR) gain increasing importance, while negative coefficients on credit variables suggest that easier financing conditions can mitigate downside risks. Notably, the influence of RECI persists throughout $h = 1$ to $h = 12$, underscoring its important role in shaping housing price dynamics across short and medium horizons.

The forecast results over the next 24 periods point to a subdued housing market, with negative year-on-year growth likely to persist and only limited prospects of recovery before mid-2027. This weak outlook is consistent with the structural break identified around April 2022, reflecting the lasting impact of demographic shifts.

From a policy perspective, several implications emerge. First, the persistence of short-term momentum suggests that immediate interventions may not quickly alter housing price dynamics, highlighting the importance of early policy signaling rather than reactive measures. Second, the medium-term role of market sentiment, inflation (CPI), and housing credit conditions (HPF) indicates that policymakers can stabilize housing prices by fostering real estate investment and easing mortgage-related financial policies. Third, the long-term influence of demographic shifts suggests that reversing housing price trends will be challenging. Therefore, policymakers should adopt a multi-horizon perspective that balances short-term stabilization needs with long-term structural considerations.

The comparison with RF shows that while machine learning methods may provide complementary forecasting strength at very long horizons, the BVS framework offers distinct advantages in transparency, interpretability, and the identification of key economic drivers. These features make BVS particularly well suited for policy environments where understanding the underlying determinants is as important as forecast accuracy. At the same time, it is important to note that this study adopts a prediction-oriented framework, focusing on forecast accuracy and horizon-specific correlations rather than establishing causal relationships. This highlights both the practical relevance of BVS for forecasting and its limitations in causal inference. Future research could address potential endogeneity concerns by incorporating instrumental variable approaches or structural VAR models, thereby complementing the predictive focus of this paper and deepening the understanding of the structural drivers of housing price dynamics.

Future research could apply this framework to housing markets in other countries, such as Korea, where local policy environments and structural conditions may significantly influence housing price dynamics. The BVS model may also be useful for forecasting regional housing markets in major metropolitan areas such as Seoul, Beijing, or Shanghai, where market fundamentals and policy responses can differ substantially from national averages. Such extensions would help assess the generalizability of the model and offer localized insights for policymakers seeking to manage housing markets under varying economic and institutional conditions.

In summary, the evidence demonstrates that BVS not only improves forecasting performance but also deepens understanding of the macro-financial forces shaping housing markets. By combining accuracy with interpretability, the framework offers a valuable tool for both research and policy, with the potential to inform housing market stability across countries and regions.

REFERENCES

Adetunji, A.B., O.N. Akande, F.A. Ajala, O. Oyewo, Y.F. Akande and G. Oluwadara (2022). "House price prediction using random forest machine learning technique," *Procedia Computer Science* 199, 806-813.

Arcelus, F. and A.H. Meltzer (1973). "The markets for housing and housing services," *Journal of Money, Credit and Banking* 5, 78-99.

Au, T.C. (2018). "Random forests, decision trees, and categorical predictors: The 'absent levels' problem," *Journal of Machine Learning Research* 19, 1-30.

Breiman, L. (2001). "Random forests," *Machine Learning* 45, 5-32.

Brown, J.P., H. Song and A. McGillivray (1997). "Forecasting UK house prices: A time varying coefficient approach," *Economic Modelling* 14, 529-548.

Brown, P.J., M. Vannucci and T. Fearn (1998). "Multivariate subspace regression with shrinkage priors," *Biometrika* 85, 635-651.

Cai, W.L., S. Chen and Y. Wang (2023). "Analysis on the influencing factors of housing price in China," *BCP Business and Management* 39, 58-64.

Case, K.E. and R.J. Shiller (1989). "The efficiency of the market for single-family homes," *American Economic Review* 79, 125-137.

Crawford, G.W. and M.C. Fratantoni (2003). "Assessing the forecasting performance of regime-switching, ARIMA and GARCH models of house prices," *Real Estate Economics* 31, 223-243.

DiPasquale, D. and W.C. Wheaton (1994). "Housing market dynamics and the future of housing prices," *Journal of Urban Economics* 35, 1-27.

Gabriel, S.A., J. Shack-Marquez and W.L. Wascher (1992). "Regional house-price dispersion and interregional migration," *Journal of Housing Economics* 2, 235-256.

Genesove, D. and C. Mayer (2001). "Loss aversion and seller behavior: Evidence from the housing market," *Quarterly Journal of Economics* 116, 1233-1260.

George, E.I. and R.E. McCulloch (1993). “Variable selection via Gibbs sampling,” *Journal of the American Statistical Association* 88, 881-889.

George, E.I. and R.E. McCulloch (1997). “Approaches for Bayesian variable selection,” *Statistica Sinica* 7, 339-373.

Gottlieb, P.D. (1976). “The effects of urban growth on the nonmetropolitan labor market,” *Review of Economics and Statistics* 58, 448-456.

Groen, J.J.J., R. Paap and F. Ravazzolo (2013). “Real-time inflation forecasting in a changing world,” *Journal of Business & Economic Statistics* 31, 29-44.

Gupta, R. and S. Das (2010). “Predicting downturns in the U.S. housing market: A Bayesian approach,” *Journal of Real Estate Finance and Economics* 41, 294-319.

Gupta, R., A. Kabundi and S.M. Miller (2011). “Forecasting the U.S. real house price index: Structural and non-structural models with and without fundamentals,” *Economic Modelling* 28, 2013-2021.

Huang, Y., D. Yi and W.A.V. Clark (2020). “Multiple home ownership in Chinese cities: An institutional and cultural perspective,” *Cities* 97, 102518.

Ishwaran, H. and J.S. Rao (2005). “Spike and slab variable selection: Frequentist and Bayesian strategies,” *Annals of Statistics* 33, 730-773.

Kang, K.H. (2018). “Mortgage loan prediction: Bayesian machine learning approach,” *Financial Stability Studies* 19, 99-129.

Koop, G. and D. Korobilis (2012). “Forecasting inflation using dynamic model averaging,” *Journal of Business & Economic Statistics* 30, 326-341.

Kuvalekar, A., S. Manchewar, S. Mahadik and S. Jawale (2020). “House price forecasting using machine learning,” *Proceedings of the 3rd International Conference on Advances in Science & Technology.*

Liang, Y. and W. Gao (2007). “Housing prices and macroeconomic factors in China: Theory and evidence from a dynamic model,” *Frontiers of Economics in China* 2, 1-23.

Maisel, S.J. (1968). "The effects of monetary policy on expenditures in specific sectors of the economy," *Journal of Political Economy* 76, 796-814.

Medeiros, M.C., G.F. Vasconcelos, Á. Veiga and E. Zilberman (2019). "Forecasting inflation in a data-rich environment: The benefits of machine learning methods," *Journal of Applied Econometrics* 34, 403-421.

Mikhed, V. and P. Zemcik (2009). "Testing for bubbles in housing markets: A panel data approach," *Journal of Real Estate Finance and Economics* 38, 366-386.

Mitchell, T.J. and J.J. Beauchamp (1988). "Bayesian variable selection in linear regression," *Journal of the American Statistical Association* 83, 1023-1032.

O'Hara, R.B. and M.J. Sillanpää (2009). "A review of Bayesian variable selection methods: What, how and which," *Bayesian Analysis* 4, 85-117.

Rapach, D.E. and J.K. Strauss (2007). "Forecasting real housing price growth in the eighth district states," *Federal Reserve Bank of St. Louis Regional Economic Development* 3, 33-42.

Rapach, D.E. and J.K. Strauss (2009). "Forecasting U.S. state-level housing price growth," *Journal of Real Estate Finance and Economics* 42, 103-130.

Satish, G.N., C.V. Raghavendran, M.D.S. Rao and C. Srinivasulu (2019). "House price prediction using machine learning," *Journal of Innovative Technology and Exploring Engineering* 8, 717-722.

Shiller, R.J. (2005). *Irrational Exuberance*, Princeton University Press.

Smith, L.B. (1969). "A model of the Canadian housing and mortgage markets," *Journal of Political Economy* 77, 795-816.

Stock, J.H. and M.W. Watson (2008). "Phillips curve inflation forecasts," NBER Working Paper No. 14322.

Teoh, E.Z., W.-C. Yau, T.S. Ong and C. Tee (2023). "Explainable housing price prediction with determinant analysis," *International Journal of Housing Markets and Analysis* 16, 1021-1045.

Thamarai, M. and S.P. Malarvizhi (2020). "House price prediction modeling using machine learning," *International Journal of Information Engineering & Electronic Business* 12, 15-20.

Truong, Q., M. Nguyen, H. Dang and B. Mei (2020). "Housing price prediction via improved machine learning techniques," *Procedia Computer Science* 174, 433-442.

Tsatsaronis, K. and H. Zhu (2004). "What drives housing price dynamics: Cross-country evidence," *BIS Quarterly Review* 3, 65-78.

Wang, X.-R., E. C.-M. Hui and J.-X. Sun (2017). "Population migration, urbanization and housing prices: Evidence from the cities in China," *Habitat International* 66, 49-56.

Xu, X. and Y. Zhang (2021). "House price forecasting with neural networks," *Intelligent Systems with Applications* 12, 200052.

Yang, Q. (2024). "Analyzing the decline of China's real estate market: Policy, market and future strategies," *Transactions on Economics, Business and Management Research* 11, 406-411.

Zhang, Y., X. Hua and L. Zhao (2012). "Exploring determinants of housing prices: A case study of Chinese experience in 1999–2010," *Economic Modelling* 29, 2349-2361.

A. APPENDIX

We first document the materials underlying Section 5.3 (Key predictors of China's housing price growth) for the short-term forecast ($h = 1$). Forecasts are produced with a Bayesian variable selection specification using the BVS2 prior and a lag length of two. The predictor set comprises 14 variables; with two lags and an intercept, the regression contains 29 coefficients in total. The Markov chain is run for 11,000 iterations in total; the first 1,000 draws are discarded as burn-in, and the remaining 10,000 draws are used for inference (MCMC size = 10,000).

We next document the materials underlying Section 5.3 (Key predictors of China's housing price growth) for the mid-term forecast ($h = 6$). Forecasts are produced with a Bayesian variable selection specification using the BVS2 prior and a lag length of two. The predictor set comprises

Parameter	Estimate	S.E.	5%	95%	Ineff
Constant	0.265	0.239	-0.128	0.657	0.781
HPI-L2	-0.879	0.050	-0.961	-0.797	0.988
RECI-L2	-0.033	0.052	-0.119	0.053	1.034
CPI-L2	0.063	0.057	-0.031	0.158	1.189
M2-L2	-0.062	0.042	-0.130	0.008	0.884
Oil-L2	-0.004	0.002	-0.008	-0.000	1.132
VAI-L2	0.001	0.017	-0.026	0.029	1.029
PMI-L2	-0.003	0.008	-0.016	0.010	0.900
CCI-L2	-0.007	0.008	-0.020	0.007	1.088
SSEC-L2	-0.003	0.003	-0.008	0.002	0.998
Bond10-L2	-0.001	0.005	-0.010	0.008	0.981
USBond10-L2	0.001	0.002	-0.002	0.004	1.089
HPF-L2	-0.008	0.011	-0.027	0.010	1.150
IHLL-L2	0.016	0.007	0.004	0.028	0.962
PNGR-L2	0.193	0.241	-0.205	0.591	1.036
HPI	1.830	0.050	1.749	1.911	0.986
RECI	0.090	0.049	0.010	0.169	1.005
CPI	-0.041	0.054	-0.129	0.050	0.997
M2	0.063	0.043	-0.008	0.134	0.989
Oil	0.003	0.002	-0.001	0.007	0.913
VAI	-0.011	0.019	-0.041	0.020	1.252
PMI	-0.007	0.007	-0.018	0.005	1.142
CCI	0.014	0.008	0.000	0.027	0.961
SSEC	0.003	0.003	-0.003	0.008	1.054
Bond10	0.002	0.006	-0.007	0.011	0.972
USBond10	-0.000	0.002	-0.003	0.003	0.753
HPF	0.007	0.012	-0.013	0.027	1.174
IHLL	-0.022	0.008	-0.035	-0.009	1.126
PNGR	-0.152	0.243	-0.556	0.242	1.119

Table A.1: Posterior summaries for $h = 1$ (BVS2, lag = 2). Ineff denotes the inefficiency factor.

14 variables; with two lags and an intercept, the regression contains 29 coefficients in total. The MCMC size = 10,000, and the first 1,000 draws are discarded as burn-in.

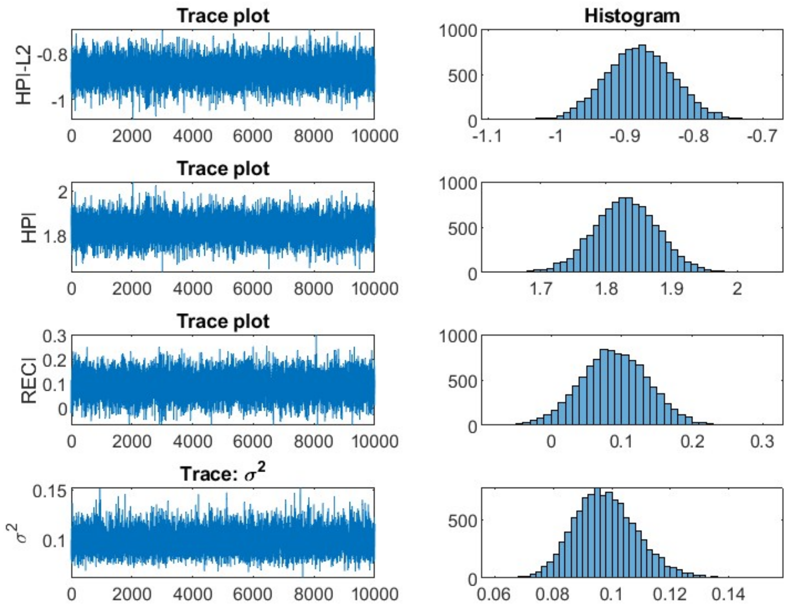

Figure A.1: MCMC convergence diagnostics for key predictors at $h = 1$. This figure shows the trace plots and posterior histograms of the estimated coefficients for HPI-L2, HPI, RECI, and the innovation variance σ^2 under the BVS2 model with lag = 2, corresponding to the one-month-ahead forecast results in Section 5.3.

We finally present the long-horizon forecast ($h = 12$) in Section 5.3. The model uses the BVS1 prior with one lag, including 14 predictors plus an intercept (15 coefficients). We run 11,000 MCMC iterations, discarding the first 1,000 as burn-in.

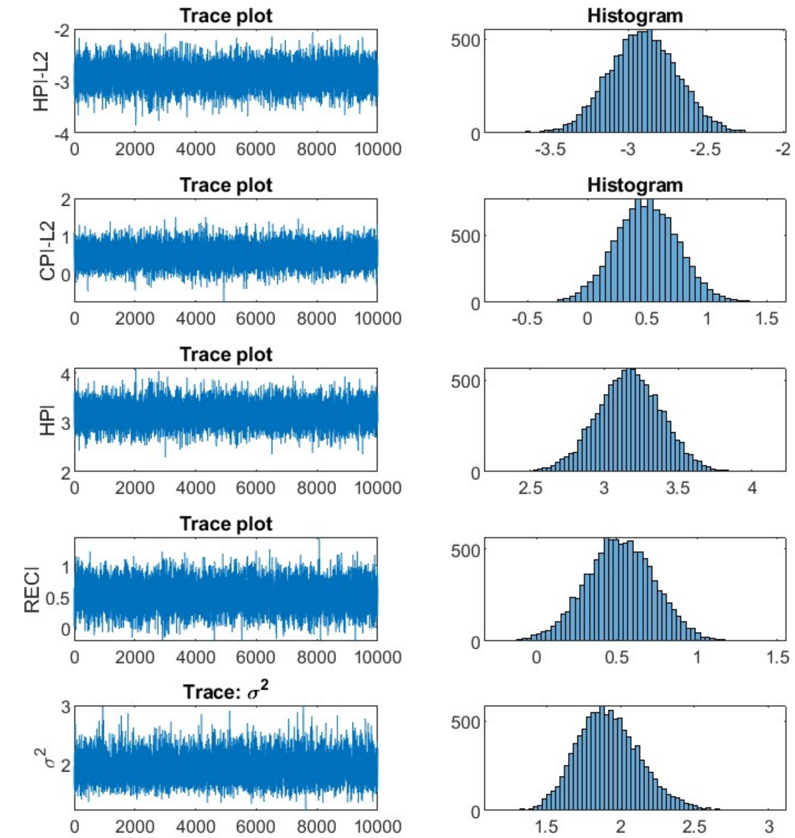

Figure A.2: MCMC convergence diagnostics for key predictors at $h = 6$. This figure shows the trace plots and posterior histograms of the estimated coefficients for HPI-L2, CPI-L2, HPI, RECI, and the innovation variance σ^2 under the BVS2 model with lag = 2, corresponding to the six-month-ahead forecast results in Section 5.3.

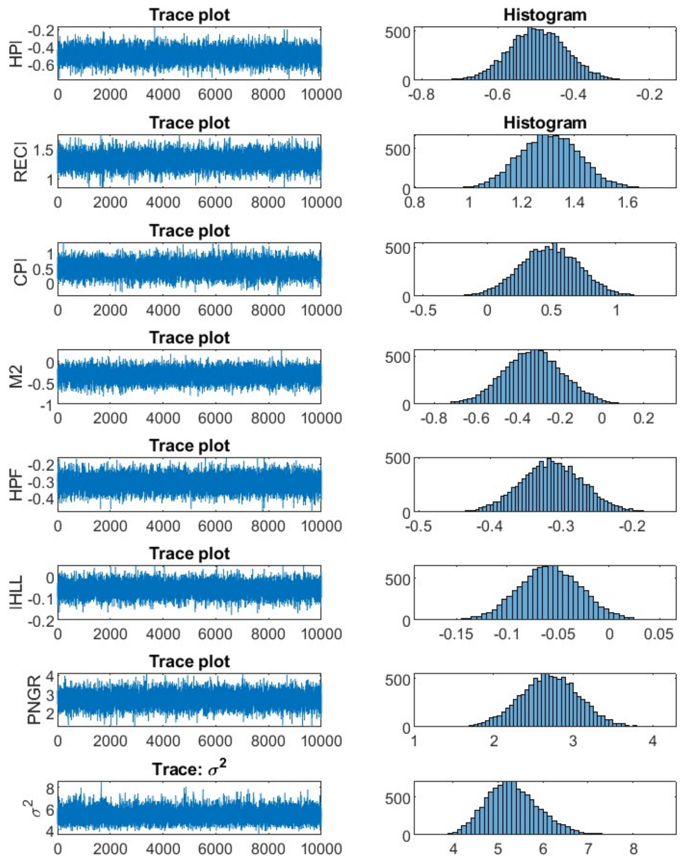

Figure A.3: MCMC convergence diagnostics for key predictors at $h = 12$. This figure shows the trace plots and posterior histograms of the estimated coefficients for HPI, RECI, CPI, M2, HPF, IHLL, PNGR, and the innovation variance σ^2 under the BVS1 model with lag = 1, corresponding to the twelve-month-ahead forecast results in Section 5.3.