

Time-Varying Stochastic Graph Learning for Probabilistic Financial Forecasting**Le Binh Trung¹, Duong Kim Ngoc¹, Le Vu Truong^{1,*}, Le Thi Kim Anh²**¹*FPT University, Ho Chi Minh City, Vietnam*²*Ho Chi Minh City University of Banking (HUB), Ho Chi Minh City, Vietnam***Corresponding author: truonglv23@fe.edu.vn*

Abstract. Financial forecasting involves both time-varying interactions among market variables and substantial uncertainty, yet these two features are often modeled separately or through fixed relational structures. This study develops a time-varying stochastic graph framework that jointly learns dynamic financial connectivity and probabilistic forecasts. The empirical application uses 2,544 daily observations from January 2015 to July 2026 for the Straits Times Index (STI), USD/SGD exchange rate, VIX, and S&P 500. A state-dependent stochastic graph is learned from the observed financial variables, and the resulting graph representations are processed by a long short-term memory network to estimate the conditional mean and standard deviation of next-day STI returns. Out-of-sample results show only marginal improvements in point forecasting over a random-walk benchmark, with directional accuracy of 50.8%. In contrast, the model produces time-varying graph connectivity and reasonably calibrated predictive intervals, with empirical coverage of 84.6% and 93.7% for the nominal 80% and 95% intervals, respectively.

1. INTRODUCTION

Forecasting financial markets remains a challenging problem because financial time series are noisy, nonlinear, and subject to changing market conditions [1–5]. The increasing availability of financial data and computational resources has motivated a transition from conventional statistical forecasting toward machine-learning and deep-learning methods capable of representing complex nonlinear patterns [6, 7]. In particular, recurrent architectures such as long short-term memory (LSTM) networks have become widely used because they can extract temporal information from sequential financial observations. Fischer and Krauss (2018) [6], for example, demonstrated the usefulness of LSTM networks for out-of-sample financial-market prediction, although they also documented that predictive profitability weakened substantially in the later part of their sample. More broadly, reviews of machine-learning applications in financial forecasting show that model

Received: Aug. 25, 2026.

2020 *Mathematics Subject Classification.* 62M20; 62M45; 91G70.*Key words and phrases.* stochastic graph; dynamic graph learning; probabilistic forecasting; financial time series.

performance depends not only on the learning algorithm but also on data representation, validation design, and the information supplied to the model.

A limitation of treating financial forecasting exclusively as a sequential learning problem is that financial variables do not evolve independently. Equity prices, exchange rates, market indices, uncertainty measures, sectors, and firms form interconnected systems in which information may propagate across related components. Financial time-series forecasting therefore involves both temporal and relational dependence [8]. Cheng et al. (2022) [9] showed that graph-based representations can combine heterogeneous financial information and account for lead-lag relationships that are difficult to represent using isolated time series. Similarly, Feng et al. (2019) [10] incorporated stock relations into temporal prediction through relational stock ranking. From a hierarchical perspective, Xu et al. (2022) [11] developed a hierarchical graph neural network for predicting the classification of price-limit-hitting stocks, further illustrating how structured relationships among stocks can be incorporated into financial prediction. These studies suggest that representing financial observations as interconnected nodes can provide information beyond that contained in their individual historical trajectories.

Graph-based financial modeling has consequently developed into an important branch of stock-market analysis. Saha et al. (2022) [12] review applications ranging from graph formulation and filtering to stock prediction and portfolio optimization, while Wang et al. (2022) [13] emphasize that graph construction itself remains a central challenge because financial graphs are rarely observed directly. Unlike applications in which network topology is naturally given, the relevant connections in financial markets must often be constructed from correlations, economic relations, external information, or learned directly from data. This distinction is important because the assumed graph determines how information is transmitted among nodes. A fixed graph constructed from historical correlations or predefined relations may summarize persistent dependence, but it implicitly assumes that the relevant relational structure remains sufficiently stable through time.

Such an assumption can be restrictive in financial applications, where relationships may evolve with market conditions. Recent studies have therefore moved from static relational representations toward dynamic graph learning. Liu et al. (2024) [14] proposed a multimodal multiscale dynamic graph convolution framework for stock-price prediction, allowing complex stock relationships to be learned jointly with information from different data sources. Qian et al. (2024) [15] similarly developed a multi-relational dynamic graph neural network in which stock relations evolve through time rather than remaining fixed. Earlier relational approaches such as Feng et al. (2019) [10] also recognized the importance of time-sensitive stock relationships. Collectively, these developments indicate that the graph itself can be treated as a dynamic object. Nevertheless, an additional modeling question remains: if the relevant relational structure is not known in advance, graph connectivity can be viewed not merely as time varying but as a latent quantity that must be learned conditional on the observed state.

A second issue concerns the forecasting target. Much of the financial machine-learning literature formulates prediction as regression, classification, ranking, or another deterministic task. Such formulations primarily emphasize the conditional location or direction of future outcomes. Financial forecasts, however, are characterized by substantial uncertainty, and two observations with similar expected returns may have very different conditional risk. This distinction has motivated recent probabilistic approaches. Daiya et al. (2024) [16], for example, proposed DiffSTOCK to combine inter-stock relational information with diffusion-based probabilistic prediction, explicitly emphasizing the difficulty of deterministic learning in financial data with a low signal-to-noise ratio. More generally, probabilistic graph-signal approaches seek to characterize a conditional distribution rather than collapse future uncertainty into a single point estimate. The recent graph-signal diffusion framework formulates financial forecasting as sampling future trajectories conditional on graph structure and observed information. These developments suggest that relational learning and uncertainty modeling should be considered jointly rather than as separate forecasting problems.

This study develops a time-varying stochastic graph probabilistic forecasting model that integrates these two dimensions within a single framework. Instead of specifying a fixed adjacency matrix, the proposed model maps each observed multivariate financial state into a matrix of edge probabilities. A differentiable stochastic relaxation generates the graph used for information propagation during training, allowing connectivity to vary across observations while remaining jointly estimable by gradient-based optimization. The resulting sequence of graph representations is processed through an LSTM temporal encoder, and the final latent representation determines both the conditional mean and conditional standard deviation of the one-step-ahead forecasting distribution. The complete system therefore jointly learns latent connectivity, temporal dependence, conditional location, and predictive uncertainty from a common likelihood objective.

The framework is evaluated using the Singapore equity market, with the Straits Times Index (STI) modeled jointly with USD/SGD exchange-rate movements, the CBOE Volatility Index (VIX), and S&P 500 returns. The application is intentionally compact: rather than constructing a large stock network, it examines whether a small heterogeneous financial system can be represented through state-dependent learned connectivity. The study makes three contributions. First, it formulates financial connectivity as a state-dependent stochastic object rather than a fixed graph, allowing the relational operator itself to change with observed market conditions. Second, it integrates this graph-learning mechanism with recurrent temporal learning and likelihood-based probabilistic forecasting, thereby estimating conditional location and uncertainty jointly. Third, the numerical application evaluates not only conventional point-forecast accuracy but also prediction-interval calibration, time variation in learned graph connectivity, and the relationship between model-implied uncertainty and subsequently realized return magnitude. In doing so, the study provides a unified applied-mathematical framework for examining dynamic relational structure and predictive uncertainty even when improvements in the conditional mean forecast are limited.

2. THEORETICAL FRAMEWORK

The proposed framework treats financial forecasting as a conditional distribution learning problem in which the relational structure among observed variables is allowed to change through time. Rather than imposing a fixed adjacency matrix, the model learns a state-dependent stochastic graph from the observed system and combines this graph representation with temporal learning to estimate the conditional distribution of the future target.

2.1. Problem formulation. Let $\mathbf{x}_t = (x_{1,t}, \dots, x_{N,t})^\top \in \mathbb{R}^N$ denote the multivariate state observed at time t . For a historical window of length L , define

$$\mathcal{X}_t = (\mathbf{x}_{t-L+1}, \dots, \mathbf{x}_t) \quad (1)$$

The objective is to estimate the one-step-ahead conditional distribution

$$p_\Theta(Y_{t+1} | \mathcal{X}_t) \quad (2)$$

rather than only a point forecast of Y_{t+1} . The model therefore seeks to represent both the expected future value and the uncertainty associated with that prediction.

The key assumption is that cross-variable dependence is not necessarily constant. We therefore define a sequence of latent graphs $G_t = (V, A_t)$, where the node set V remains fixed, but the adjacency matrix A_t varies with the observed state. In contrast to a static graph satisfying $A_t = A$ for all t , the proposed framework permits the relational structure itself to evolve over time.

2.2. Time-varying stochastic graph. At each time t , a nonlinear mapping transforms the observed state into an edge-logit matrix, $L_t = f_{\theta_G}(\mathbf{x}_t)$ from which the edge-probability matrix is obtained as

$$P_t = \sigma(L_t) \quad (3)$$

Thus, $P_t = P_{\theta_G}(\mathbf{x}_t)$ and the graph connectivity is explicitly state-dependent. To retain stochastic graph formation while preserving differentiability, the model uses a continuous logistic relaxation. For $u_{ij,t} \sim U(0, 1)$, the relaxed stochastic edge is defined as

$$a_{ij,t} = \sigma\left(\frac{\ell_{ij,t} + \log u_{ij,t} - \log(1 - u_{ij,t})}{\tau}\right) \quad (4)$$

where $\tau > 0$ is a temperature parameter. The resulting matrix $A_t = [a_{ij,t}]$ is a stochastic continuous approximation to a binary adjacency matrix. During model evaluation, the stochastic perturbation is removed, and the conditional edge probabilities P_t are used directly.

The graph is row-normalized and applied to the observed state. The normalized adjacency matrix is $\tilde{A}_t = D_t^{-1}A_t$ and the corresponding graph representation is

$$\mathbf{h}_t = g_{\theta_H}(\tilde{A}_t \mathbf{x}_t) \quad (5)$$

where D_t is the row-degree matrix and $g_{\theta_H}(\cdot)$ is a nonlinear transformation. Equation (5) shows that the operator acting on $\mathbf{x}_t$ is itself state-dependent. The transformation is therefore not equivalent to applying a fixed linear weighting matrix to the variables.

Under the adopted matrix convention,

$$a_{ij,t} : j \rightarrow i \quad (6)$$

so that column j represents a source node and row i a receiving node. Applying this transformation over the complete historical window produces the sequence

$$H_t = (\mathbf{h}_{t-L+1}, \dots, \mathbf{h}_t) \quad (7)$$

which represents the evolution of the learned graph states through time.

2.3. Temporal probabilistic forecasting. The sequence H_t is processed by a recurrent temporal mapping,

$$\mathbf{z}_t = \mathcal{R}_{\theta_R}(H_t) \quad (8)$$

where $\mathcal{R}_{\theta_R}$ is implemented as a long short-term memory network. The recurrent representation $\mathbf{z}_t$ summarizes the temporal evolution of the graph-transformed states over the historical window. The location parameter of the predictive distribution is generated as $\mu_{t+1} = f_{\theta_\mu}(\mathbf{z}_t)$. The scale parameter is defined as

$$\sigma_{t+1} = \text{Softplus}(f_{\theta_\sigma}(\mathbf{z}_t)) + \delta \quad (9)$$

where $\delta > 0$ guarantees a strictly positive scale. The conditional forecast is specified as

$$Y_{t+1} | \mathcal{X}_t \sim \mathcal{N}(\mu_{t+1}, \sigma_{t+1}^2) \quad (10)$$

Thus, μ_{t+1} represents the conditional location, whereas σ_{t+1} represents model-implied predictive uncertainty. All graph, temporal, and output parameters are estimated jointly by minimizing the Gaussian negative log-likelihood,

$$\Theta^* = \arg \min_{\Theta} \frac{1}{M} \sum_{t=1}^M \left[\log \sigma_{t+1} + \frac{(y_{t+1} - \mu_{t+1})^2}{2\sigma_{t+1}^2} \right] \quad (11)$$

up to the constant Gaussian term. The parameter set can be written as $\Theta = \{\theta_G, \theta_H, \theta_R, \theta_\mu, \theta_\sigma\}$. Consequently, the latent graph is not estimated independently and then supplied to a separate forecasting model. Instead, the entire sequence $\mathcal{X}_t \rightarrow P_t \rightarrow H_t \rightarrow \mathbf{z}_t \rightarrow (\mu_{t+1}, \sigma_{t+1})$ is optimized jointly for probabilistic forecasting. The learned graph should therefore be interpreted as a latent predictive connectivity structure rather than as a causal network. Its edge probabilities identify the relational structure that the model finds useful for representing the conditional distribution of the target, but they do not constitute structural or causal coefficients.

The proposed framework is particularly applicable to financial systems, where observed variables are interconnected, and their dependence structure may change with market conditions. Graph-based representations have increasingly been used in stock-market analysis because financial assets can naturally be represented as nodes connected through statistical, economic, or

learned relationships [12]. More recent studies have moved beyond fixed financial networks toward dynamic relational learning. For example, Liu et al. (2024) [14] developed a dynamic graph convolution framework to capture evolving relationships in stock-price prediction, while Qian et al. (2024) [15] proposed a multi-relational dynamic graph neural network to accommodate the changing and heterogeneous interactions underlying stock movements. Probabilistic relational modeling has also emerged as an important extension: Daiya et al. (2024) [16] combined inter-stock relationships with diffusion-based probabilistic forecasting, explicitly motivated by the high uncertainty and low signal-to-noise ratio of financial data. These developments support the use of graph-based models when both temporal dynamics and relational dependence matter. The present framework follows this general direction but models the relational structure as a state-dependent stochastic graph, $P_t = P_{\theta_G}(\mathbf{x}_t)$, and jointly maps the resulting graph sequence into the conditional location and scale of the forecasting distribution. This formulation is suitable not only for relationships among individual stocks, but also for systems containing different classes of financial variables—such as equity returns, exchange-rate movements, global market conditions, and uncertainty measures—whose predictive connections may evolve over time.

3. METHODOLOGY

3.1. Data and Variables. The empirical application examines the Singapore equity market using daily financial data from January 2015 to July 31, 2026. Four market variables are considered: the Straits Times Index (STI), the USD/SGD exchange rate, the CBOE Volatility Index (VIX), and the S&P 500 index. The series were obtained from Yahoo Finance using their corresponding market tickers. The selection provides a compact multivariate system combining the domestic equity market with exchange-rate movements, global market uncertainty, and international equity-market conditions. The STI represents the domestic market and provides the forecasting target. The USD/SGD exchange rate is included to capture movements in the external value of the Singapore dollar, which may be relevant to a highly internationally integrated equity market. The S&P 500 represents broad global, and particularly U.S., equity-market conditions. The VIX provides a complementary measure of market uncertainty. Unlike the three price-based variables, which are transformed into continuously compounded returns, the VIX is retained in levels because its level directly represents the prevailing state of market uncertainty.

Let P_t^{STI} , P_t^{SGD} , and P_t^{SP500} denote the closing values of the STI, USD/SGD exchange rate, and S&P 500, respectively. Their continuously compounded daily returns are calculated as

$$r_t^k = \ln \left(\frac{P_t^k}{P_{t-1}^k} \right) \quad (12)$$

where $k \in \{STI, SGD, SP500\}$. Let VIX_t denote the observed VIX level. The four-dimensional state vector supplied to the graph model is therefore

$$\mathbf{x}_t = \left(r_t^{STI}, r_t^{SGD}, VIX_t, r_t^{SP500} \right)^\top \quad (13)$$

The forecasting target is the next-day STI return,

$$Y_{t+1} = r_{t+1}^{STI} \quad (14)$$

so that information available through date t is used exclusively to forecast the STI return at $t+1$. The variables in the study are described in Table 1.

TABLE 1. Definition and construction of variables

Variable	Symbol	Construction	Role in the model
STI return	r_t^{STI}	$\ln(STI_t / STI_{t-1})$	Domestic equity-market node; historical values also define the forecasting target
SGD return	r_t^{SGD}	$\ln(SGD_t / SGD_{t-1})$	Exchange-rate node representing daily USD/SGD movements
VIX level	VIX_t	Daily closing level	Global market-uncertainty node
S&P 500 return	r_t^{SP500}	$\ln(SP500_t / SP500_{t-1})$	Global equity-market node
Future STI return	Y_{t+1}	r_{t+1}^{STI}	One-day-ahead forecasting target

After transformation, observations were aligned by trading date, and incomplete observations were removed, resulting in 2,544 common daily observations. Table 1 summarizes the definition, construction, and modeling role of each variable. Importantly, the four variables are treated as graph nodes rather than as predetermined explanatory variables with fixed coefficients. Their effective contribution to the forecasting system is determined endogenously through the time-varying stochastic connectivity mechanism introduced in Section 2.

Before model estimation, each input variable was standardized using the mean and standard deviation estimated from the training sample only. The STI forecasting target was standardized separately using its corresponding training-sample parameters. The same transformations were subsequently applied to the validation and test samples without re-estimation. This procedure preserves the chronological forecasting structure and prevents information from future observations from entering data preprocessing.

3.2. Forecasting Model. The empirical forecasting model implements the theoretical framework described in Section 2 using a time-varying stochastic graph combined with an LSTM temporal encoder and a Gaussian probabilistic output layer. The model takes a rolling sequence of the four standardized financial variables as input and produces the conditional mean and conditional standard deviation of the next-day STI return.

Input sequence and graph construction. Each forecasting observation is constructed using a rolling window of 20 trading days. Accordingly, the model input at time t is

$$\mathcal{X}_t = (\mathbf{x}_{t-19}, \dots, \mathbf{x}_t) \quad (15)$$

where each $\mathbf{x}_t$ contains the standardized STI return, USD/SGD return, VIX level, and S&P 500 return. The resulting input tensor, therefore, contains 20 consecutive observations for four financial nodes. The graph-learning module processes each daily state separately. A feed-forward neural network maps the four-dimensional state into a 4×4 matrix of edge logits. The graph network contains an intermediate hidden layer with a ReLU activation and produces 16 outputs, which are reshaped into the adjacency structure. Self-connections are excluded so that the graph captures only interactions between distinct financial variables. During training, stochastic edge weights are generated from the learned logits using the logistic relaxation introduced in Section 2. The temperature parameter is fixed at $\tau = 0.5$. The resulting adjacency matrix is normalized row-wise before graph propagation. In inference, stochastic perturbations are removed, and the sigmoid-transformed edge probabilities are used directly. This provides a deterministic conditional graph for each observed market state while preserving stochastic graph learning during parameter estimation. For each day in the 20-day window, the normalized graph aggregates information across the four financial nodes. The aggregated values are subsequently passed through a nonlinear graph-transformation layer. Repeating this operation across the complete input window generates a sequence of 20 graph-transformed financial states.

Temporal encoding and probabilistic output. The sequence of graph representations is supplied to an LSTM network to capture temporal dependence across successive market states. The LSTM uses a hidden dimension of 64, and its final hidden representation summarizes information from the complete 20-day historical window. The final LSTM representation is passed to two separate output heads. The first generates the conditional mean of the standardized next-day STI return, while the second generates its conditional standard deviation. A Softplus transformation is applied to the scale output to ensure positivity [17, 18].

The model, therefore, produces $(\widehat{\mu}_{t+1}, \widehat{\sigma}_{t+1})$ for every forecasting observation. These parameters define the estimated Gaussian predictive distribution

$$Y_{t+1} | \mathcal{X}_t \sim \mathcal{N}(\widehat{\mu}_{t+1}, \widehat{\sigma}_{t+1}^2) \quad (16)$$

For out-of-sample evaluation, the predictive mean and standard deviation are transformed back to the original STI-return scale for the calculation of MAE, RMSE, directional accuracy, and prediction intervals. Gaussian negative log-likelihood is evaluated on the standardized target scale used during model estimation. The predictive mean provides the conventional point forecast, whereas the estimated standard deviation determines the width of the probabilistic prediction intervals. The four financial variables first enter the state-dependent stochastic graph module. Graph propagation is performed separately for each observation in the 20-day input window, after which the resulting sequence is processed by the LSTM temporal encoder. The final recurrent

representation is mapped into the conditional mean and conditional standard deviation of the next-day STI return.

Model estimation. The model is estimated by minimizing the Gaussian negative log-likelihood defined in Section 2. All components—including the edge-generation network, graph-transformation network, LSTM encoder, and distributional output heads—are trained jointly through back-propagation. Optimization is performed using the Adam algorithm with an initial learning rate of 0.001 and a weight decay of 0.00001. Gradient norms are clipped at 1.0 to reduce the effect of unstable parameter updates arising during stochastic graph learning. Training is conducted for a maximum of 300 epochs with early stopping based on validation negative log-likelihood. The parameter configuration producing the lowest validation loss is retained, and training is terminated when no further improvement is observed for 30 consecutive epochs.

The final implementation can therefore be summarized as a sequential computational pipeline:

$$\begin{aligned} 20\text{-day financial states} &\rightarrow \text{time-varying stochastic graphs} \rightarrow \text{graph representations} \\ &\rightarrow \text{LSTM} \rightarrow (\widehat{\mu}_{t+1}, \widehat{\sigma}_{t+1}) \end{aligned}$$

This specification preserves a distinction between cross-variable and temporal information. The graph module determines how the four financial variables interact at each date, whereas the LSTM learns how these graph-transformed states evolve across the 20-day historical window. The Gaussian output layer then converts the resulting representation into a probabilistic one-day-ahead forecast.

3.3. Model Evaluation. The forecasting experiment follows a chronological design to preserve the temporal ordering of the data. Observations from 2015–2023 are used for model training, 2024 is reserved for validation and model selection, and the period from January 2025 to July 2026 is used exclusively for out-of-sample evaluation. After constructing the 20-day input sequences, the resulting samples contain 1,954 training, 199 validation, and 331 test observations. Because input sequences were constructed separately within the training, validation, and test partitions, the first 20 observations associated with each partition were not used as forecasting samples, resulting in 60 fewer forecasting observations than the 2,544 aligned daily observations. Point-forecast performance is evaluated using mean absolute error (MAE), root mean squared error (RMSE), and directional accuracy (DA) [19,20]. The predictive mean of the proposed model is compared with a random-walk benchmark that assumes a zero expected next-day return, $\widehat{r}_{t+1}^{RW} = 0$.

Because the model produces a complete conditional distribution, evaluation also considers Gaussian negative log-likelihood (NLL) and the empirical coverage rates of the 80% and 95% prediction intervals. These measures assess whether the estimated predictive distribution provides an appropriate representation of forecast uncertainty in addition to point accuracy. Finally, the information contained in the estimated conditional standard deviation is examined by comparing $\widehat{\sigma}_{t+1}$ with the magnitude of the subsequently realized return, $|r_{t+1}^{STI}|$. Both Pearson and Spearman

correlations are reported. Pearson correlation measures linear association, while Spearman correlation provides a rank-based assessment that is less sensitive to extreme observations. The learned edge probabilities are additionally examined over the test period to evaluate whether the estimated graph connectivity exhibits meaningful time variation.

4. RESULTS

4.1. Descriptive Statistics and Correlation Analysis. Table 2 reports the descriptive statistics for the four financial variables over the full sample. The three return series have daily means close to zero. The STI return has a mean of 0.000208 and a standard deviation of 0.008440, while the S&P 500 return exhibits greater variability, with a standard deviation of 0.011361. USD/SGD movements are comparatively smaller, with a standard deviation of 0.003100. The VIX has a mean of 18.42 and a median of 16.76, but reaches a maximum of 82.69, reflecting the presence of episodes of exceptionally high market uncertainty.

TABLE 2. Descriptive statistics

Variable	Mean	SD	Min	Q1	Median	Q3	Max
STI return	0.000208	0.00844	-0.07752	-0.00381	0.000275	0.004549	0.058946
USD/SGD return	0.00000	0.0031	-0.01629	-0.00193	0.00000	0.001752	0.014956
VIX level	18.42427	7.134537	9.14	13.705	16.76	21.1325	82.69
S&P 500 return	0.000417	0.011361	-0.12765	-0.00379	0.000733	0.005796	0.090895

The ranges of the return series also indicate the occurrence of substantial market movements. In particular, the minimum daily returns are -7.75% for the STI and -12.77% for the S&P 500, considerably larger in magnitude than their respective standard deviations. This feature is relevant to the subsequent probabilistic analysis because forecast uncertainty may vary considerably across normal and extreme market conditions. Table 3 presents the Pearson correlation matrix. The unconditional linear relationships among the variables are generally modest. STI returns have a positive correlation of 0.230 with S&P 500 returns and negative correlations of -0.202 with USD/SGD returns and -0.164 with the VIX. The correlation between S&P 500 returns and the VIX is also negative at -0.187. The remaining pairwise correlations are close to zero.

TABLE 3. Pearson correlation matrix

Variable	STI return	USD/SGD return	VIX level	S&P 500 return
STI return	1	-0.202	-0.164	0.23
USD/SGD return	-0.202	1	0.048	-0.014
VIX level	-0.164	0.048	1	-0.187
S&P 500 return	0.23	-0.014	-0.187	1

The correlation matrix does not indicate a uniformly strong static dependence structure among the four variables. However, these coefficients summarize unconditional linear associations over

the entire sample and do not capture possible changes in relational structure across market states. Whether the forecasting model identifies such time variation is examined subsequently through the learned graph probabilities.

4.2. Forecasting Performance and Predictive Distribution. Table 4 reports the out-of-sample forecasting performance over the test period. The time-varying stochastic graph model produces an MAE of 0.0054 and an RMSE of 0.0086, compared with 0.0054 and 0.0087, respectively, for the random-walk benchmark. The improvement in point-forecast accuracy is therefore small. Directional accuracy is 0.5076, indicating that the model correctly predicts the sign of the next-day STI return in approximately 50.8% of the test observations.

TABLE 4. Out-of-sample forecasting performance

Model	MAE	RMSE	Direction accuracy	Gaussian NLL	80% coverage	95% coverage
TV Stochastic Graph	0.0054	0.0086	0.5076	1.3679	0.8459	0.9366
Random Walk	0.0054	0.0087	–	–	–	–

Note. NLL denotes Gaussian negative log-likelihood. Coverage denotes the proportion of realized STI returns falling within the corresponding model-implied prediction intervals. MAE, RMSE, directional accuracy, and interval coverage are reported on the original STI-return scale, whereas Gaussian NLL is computed on the standardized target scale.

The limited difference relative to the random-walk benchmark indicates that the proposed model does not generate a substantial improvement in the conditional mean forecast. This pattern is also visible in Fig. 1, where the predictive mean remains close to zero during most of the test period despite considerable variation in realized STI returns.

The probabilistic results provide additional information beyond point accuracy. The test Gaussian NLL is 1.3679. The nominal 80% prediction interval achieves an empirical coverage rate of 84.59%, while the nominal 95% interval achieves a coverage rate of 93.66%. The former therefore exhibits moderate overcoverage, whereas the latter shows slight undercoverage. Nevertheless, both empirical rates remain reasonably close to their nominal levels. Fig. 1 further shows that the width of the predictive distribution is not constant over time. The most pronounced expansion occurs around April 2025, coinciding with exceptionally large movements in realized STI returns. The prediction intervals subsequently contract as market movements return toward more typical levels and widen again, although less sharply, during several later periods. These results indicate that the distributional component of the model responds to changing market conditions even though the conditional mean remains difficult to predict.

The out-of-sample evidence distinguishes between two dimensions of forecasting performance. The model provides only marginal gains in point prediction relative to the random-walk benchmark, whereas its estimated predictive distribution exhibits time-varying dispersion and empirical




FIGURE 1. Out-of-sample probabilistic forecasts of daily STI returns.

Note. The figure presents realized STI returns, the model-implied predictive mean, and the 80% and 95% prediction intervals over the test period from January 2025 to July 2026.

coverage reasonably close to the intended probability levels. The behavior of this conditional uncertainty is examined more directly in the subsequent analysis.

4.3. Time-Varying Graph Connectivity. Table 5 reports the average learned edge probabilities over the out-of-sample period. Under the graph convention adopted in the model, columns represent source nodes and rows represent receiving nodes. The estimated probabilities, therefore, describe the model-implied predictive connectivity between pairs of financial variables rather than conventional correlation coefficients.

TABLE 5. Average learned edge probabilities during the out-of-sample period

Target \ Source	STI	USD/SGD	VIX	S&P 500
STI	–	0.41	0.633	0.374
USD/SGD	0.409	–	0.569	0.46
VIX	0.636	0.374	–	0.425
S&P 500	0.693	0.423	0.287	–

Note. rows = target, columns = source

For the STI node, the strongest average incoming connection originates from the VIX, with an edge probability of 0.633. The corresponding probabilities for USD/SGD and the S&P 500 are 0.410 and 0.374, respectively. Thus, among the three external nodes considered in the model, global market uncertainty receives the largest average relational weight in constructing the graph representation of the STI. The learned graph is also asymmetric. For example, the average S&P 500-to-STI probability is 0.374, whereas the reverse STI-to-S&P 500 probability is considerably larger at

0.693. Similarly, the VIX-to-STI and STI-to-VIX probabilities are 0.633 and 0.636, respectively. The graph therefore does not impose reciprocal relationships between nodes; each directed connection is learned separately as part of the forecasting system. More importantly, the average probabilities in Table 5 conceal substantial variation through time. Fig. 2 presents the three learned incoming connections to the STI over the test period. The VIX-to-STI edge remains the strongest for most observations, generally fluctuating at relatively high probability levels. In contrast, the USD/SGD-to-STI and S&P 500-to-STI connections exhibit lower average levels and more pronounced changes during several periods.




FIGURE 2. Time-varying learned edge probabilities into the STI node during the out-of-sample period.

Note. The figure reports the model-implied edge probabilities from the VIX, USD/SGD, and S&P 500 nodes to the STI node from January 2025 to July 2026.

The temporal movements shown in Fig. 2 provide direct numerical evidence that the fitted relational structure is not constant over the out-of-sample period. In particular, the relative importance assigned to the three incoming nodes changes across market states rather than being determined by a single fixed adjacency matrix. This behavior is consistent with the state-dependent construction of the graph, in which the edge-probability matrix is generated as a function of the contemporaneous financial state. These probabilities should nevertheless be interpreted as learned predictive connectivity, rather than as causal transmission coefficients. Because the graph is estimated jointly with the forecasting objective, a larger edge probability indicates that the corresponding connection receives greater weight within the model's latent representation. It does not establish that changes in the source variable causally affect the receiving variable.

4.4. Conditional Uncertainty. The probabilistic specification allows the model to generate a time-varying conditional standard deviation in addition to the predictive mean. Fig. 3 compares the

estimated conditional standard deviation with the absolute realized STI return over the out-of-sample period. The predicted uncertainty is clearly nonconstant and increases substantially during periods characterized by large market movements.




FIGURE 3. Predicted conditional uncertainty and absolute realized STI returns over the out-of-sample period.

Note. The figure compares the model-implied conditional standard deviation with the absolute realized daily STI return from January 2025 to July 2026.

The most pronounced increase in predicted uncertainty occurs around April 2025, when the STI experiences unusually large daily movements. The estimated standard deviation subsequently declines toward its lower range as market conditions stabilize. Smaller increases are also observed during later episodes of elevated realized return magnitude. These dynamics indicate that the scale component of the predictive distribution responds to changes in market conditions rather than remaining approximately constant throughout the test period.

The relationship is examined more directly in Fig. 4, which plots the estimated conditional standard deviation against the absolute subsequently realized STI return. The Pearson correlation between the two quantities is positive at 0.3295, indicating that observations assigned greater predictive uncertainty tend to be associated with larger realized return magnitudes.

The rank-based association is weaker. The Spearman correlation is 0.119 and remains statistically significant at the 5% level ($p=0.0307$). The difference between the Pearson and Spearman coefficients indicates that the stronger linear association is partly influenced by observations associated with relatively large market movements. Thus, the estimated conditional standard deviation does not provide a strong ranking of day-to-day realized return magnitudes across the entire test sample. These results suggest that the model captures a meaningful but imperfect form of time-varying predictive uncertainty. The conditional scale responds particularly to episodes of elevated market volatility, while its association with ordinary day-to-day variations in return magnitude




FIGURE 4. Association between predicted conditional standard deviation and absolute realized STI return.

Note. The solid line represents the least-squares linear fit. The Pearson correlation between predicted conditional standard deviation and absolute realized return is 0.3295.

remains relatively weak. Accordingly, $\widehat{\sigma}_{t+1}$ is best interpreted as a model-implied measure of predictive uncertainty rather than as a precise forecast of subsequent absolute returns.

5. DISCUSSION

The empirical results provide a differentiated assessment of the proposed time-varying stochastic graph framework. The model does not substantially improve the conditional mean forecast of daily STI returns relative to a simple random-walk benchmark. However, it learns a clearly time-varying relational structure among the financial variables and generates a nonconstant conditional scale that responds to periods of elevated market movements. These findings suggest that the main value of the framework lies in jointly representing dynamic connectivity and predictive uncertainty rather than in extracting a strong directional signal from daily returns.

5.1. Dynamic financial connectivity. The learned graph provides evidence that financial connectivity is neither uniform across variables nor constant through time. Among the incoming connections to the STI, the VIX receives the highest average edge probability, followed by USD/SGD and the S&P 500. More importantly, Fig. 2 shows that these probabilities evolve throughout the out-of-sample period. The fitted graph, therefore, differs fundamentally from a representation based on a single correlation or adjacency matrix.

This result is consistent with the growing use of dynamic relational structures in financial prediction. While earlier graph-based approaches demonstrated the usefulness of incorporating stock relationships into predictive models [9–11], more recent models explicitly allow these relationships to evolve over time [14,15]. The present results support this general direction. In particular, the relatively modest unconditional correlations reported in Table 3 coexist with considerably stronger learned edge probabilities and visible temporal changes in those probabilities. This difference is not contradictory because the two quantities represent different objects. Pearson correlation summarizes average linear co-movement over the sample, whereas the graph probabilities are nonlinear, state-dependent quantities learned jointly with the forecasting objective.

The strong average VIX-to-STI connection is particularly notable. It indicates that the model consistently assigns substantial predictive weight to global market uncertainty when constructing the latent representation of the Singapore equity market. However, this result should not be interpreted as evidence that changes in the VIX causally determine STI returns. The graph is optimized to improve representation of the forecasting distribution rather than to identify structural economic relationships. Accordingly, its edge probabilities are better understood as measures of latent predictive connectivity.

The asymmetric structure of the estimated graph reinforces this interpretation. Several connections differ considerably depending on their direction, most notably the STI–S&P 500 pair. Because each directed edge is estimated independently, the model does not impose reciprocal relationships between financial variables. Such flexibility is useful in financial systems where the predictive information associated with one variable need not be symmetric with the information flowing in the opposite direction.

5.2. Forecastability and predictive uncertainty. A second important finding is the contrast between conditional-mean predictability and conditional uncertainty. The proposed model produces only marginal reductions in MAE and RMSE relative to the random-walk benchmark, while directional accuracy remains close to 50%. The flexibility introduced by the stochastic graph and recurrent architecture therefore does not translate into substantial one-day-ahead predictability of the direction or conditional mean of STI returns.

This negative result is informative rather than inconsistent with the proposed framework. Financial return series have low signal-to-noise ratios, and flexible nonlinear architectures cannot create predictable information when the conditional mean contains little exploitable signal. The result is also consistent with the broader observation that strong in-sample representation capacity does not necessarily imply economically meaningful out-of-sample return predictability [6,7]. In the present application, the graph should therefore not be interpreted as a mechanism that makes daily equity returns strongly predictable. The probabilistic component provides a different picture. Empirical coverage rates of the 80% and 95% prediction intervals are reasonably close to their nominal levels, while the estimated conditional standard deviation changes considerably across the test period. In particular, the model produces a sharp increase in predictive uncertainty during

the large market movements observed around April 2025. The scale parameter, therefore, captures information that is largely absent from the conditional mean.

The association between predicted uncertainty and realized return magnitude nevertheless requires a calibrated interpretation. The Pearson correlation of 0.3295 indicates a moderate positive linear relationship, whereas the substantially smaller Spearman correlation of 0.119 indicates only a weak monotonic association. The difference suggests that part of the Pearson relationship is generated by relatively extreme market episodes. The model appears better at recognizing periods in which uncertainty becomes unusually elevated than at precisely ranking the magnitude of ordinary day-to-day returns. Consequently, the estimated conditional standard deviation should not be viewed as a direct forecast of future absolute returns. Rather, it represents the dispersion of the model-implied predictive distribution. The results suggest that this dispersion contains meaningful information about changing market conditions, but its correspondence with subsequently realized volatility is imperfect. This distinction is important because probabilistic forecasting concerns the calibration and dispersion of the entire predictive distribution, not simply the prediction of a volatility proxy.

The findings are therefore consistent with the motivation for probabilistic relational forecasting. Daiya et al. (2024) [16], for example, emphasized the value of incorporating relational information into a probabilistic framework rather than reducing financial prediction to deterministic outcomes. The present results provide a related implication from a substantially simpler distributional specification: even when improvements in point forecasts are negligible, jointly estimated relational structure and conditional uncertainty can remain informative characteristics of the forecasting system.

5.3. Implications of the proposed framework. From a modeling perspective, the main contribution of the framework is the treatment of graph connectivity as an endogenous state-dependent object. Conventional graph forecasting requires either an externally specified network or a graph constructed from historical relationships. The present formulation instead defines $P_t = P_{\theta_G}(\mathbf{x}_t)$ so that the relational operator itself changes with the observed system. Combined with stochastic relaxation during training, this allows uncertainty in relational structure to enter the learning process while preserving end-to-end differentiability. This formulation separates three objects that are often embedded implicitly within a single forecasting function: cross-variable connectivity, temporal dependence, and predictive uncertainty. The graph module determines how information is combined across variables at each point in time; the recurrent component summarizes the evolution of these graph states; and the probabilistic output layer maps the resulting representation into conditional location and scale. All three components are optimized through a common likelihood objective.

The empirical application also demonstrates that a graph-learning framework need not require a large network of individual stocks. The four-node system combines domestic equity returns, exchange-rate movements, global equity returns, and market uncertainty. This heterogeneous

construction is particularly relevant when the forecasting problem concerns interactions among different classes of financial variables rather than relationships among securities of the same type. It also complements evidence that equity-market modeling can benefit from considering exchange-rate exposure and other cross-market information [8].

At the same time, the results establish clear boundaries for interpretation. Learned connectivity does not imply causal transmission, and architectural complexity does not guarantee superior point forecasts. The empirical value of the proposed framework is that it provides a unified mathematical representation in which the relational structure and forecasting distribution can both adapt to the observed state. This feature makes the framework potentially applicable beyond the specific STI experiment to other multivariate stochastic systems characterized by evolving interactions and heterogeneous uncertainty.

Several extensions follow naturally from these findings. The Gaussian specification could be replaced by heavier-tailed or asymmetric conditional distributions to accommodate extreme financial returns more explicitly. Larger graphs could incorporate additional international markets, currencies, commodities, or macro-financial variables. Alternative graph regularization schemes could also be introduced to encourage sparsity or persistence in the learned topology. Finally, multi-step forecasting would allow the persistence of both dynamic connectivity and predictive uncertainty to be examined over longer horizons.

6. CONCLUSION

This study developed a time-varying stochastic graph framework for probabilistic financial forecasting and applied it to the Singapore equity market using the STI, USD/SGD, VIX, and S&P 500. The model jointly learns state-dependent graph connectivity, temporal dependence, and the conditional distribution of next-day STI returns. The empirical results show only marginal improvement in point forecasting relative to the random-walk benchmark, with directional accuracy remaining close to 50%. However, the learned graph exhibits clear time variation, with the VIX providing the strongest average incoming connection to the STI. The probabilistic component also produces prediction-interval coverage reasonably close to nominal levels and a time-varying conditional scale that responds to episodes of elevated market movements. Overall, the findings suggest that the main value of the framework lies in representing dynamic financial connectivity and predictive uncertainty rather than substantially improving short-horizon conditional-mean forecasts.

Several limitations provide directions for further research. The empirical application uses a relatively small four-node financial system, assumes a Gaussian conditional distribution, and considers only one-day-ahead forecasting for a single equity market. Moreover, the learned edge probabilities represent predictive connectivity and should not be interpreted as causal relationships. Future research could extend the framework to larger financial networks and multiple markets, adopt heavy-tailed or asymmetric predictive distributions, introduce alternative graph

regularization mechanisms, and examine multi-step forecasting. Such extensions may provide a richer representation of both evolving financial connectivity and uncertainty under more complex market conditions.

Conflicts of Interest: The authors declare that there are no conflicts of interest regarding the publication of this paper.

REFERENCES

- [1] D.P. Gandhmal, K. Kumar, Systematic Analysis and Review of Stock Market Prediction Techniques, *Comput. Sci. Rev.* 34 (2019), 100190. <https://doi.org/10.1016/j.cosrev.2019.08.001>.
- [2] H. Rezaei, H. Faaljou, G. Mansourfar, Stock Price Prediction Using Deep Learning and Frequency Decomposition, *Expert Syst. Appl.* 169 (2021), 114332. <https://doi.org/10.1016/j.eswa.2020.114332>.
- [3] C. Zhang, Q. Lan, X. Mi, Z. Zhou, C. Ma, et al., A Denoising Method Based on the Nonlinear Relationship between the Target Variable and Input Features, *Expert Syst. Appl.* 218 (2023), 119585. <https://doi.org/10.1016/j.eswa.2023.119585>.
- [4] H. Yu, J. Dong, B. Li, Economic Insight through an Optimized Deep Learning Model for Stock Market Prediction Regarding Dow Jones Industrial Average Index, *Expert Syst. Appl.* 291 (2025), 128473. <https://doi.org/10.1016/j.eswa.2025.128473>.
- [5] Y. Huang, C. Yang, Enhancing Stock Price Forecasting with a Modular Deep Learning Framework Incorporating Plug-and-Play Transformer Variants, *Expert Syst. Appl.* 315 (2026), 131572. <https://doi.org/10.1016/j.eswa.2026.131572>.
- [6] T. Fischer, C. Krauss, Deep Learning with Long Short-Term Memory Networks for Financial Market Predictions, *Eur. J. Oper. Res.* 270 (2018), 654–669. <https://doi.org/10.1016/j.ejor.2017.11.054>.
- [7] K. Olorunnimbe, H. Viktor, Deep Learning in the Stock Market—a Systematic Survey of Practice, Backtesting, and Applications, *Artif. Intell. Rev.* 56 (2023), 2057–2109. <https://doi.org/10.1007/s10462-022-10226-0>.
- [8] T.T. Huynh, B.T. Khoa, N.T.K. An, Return Reversal in Portfolios Optimized under Exchange Rate Risk: Evidence from Vietnam’s HNX Market, *Int. J. Anal. Appl.* 23 (2025), 237. <https://doi.org/10.28924/2291-8639-23-2025-237>.
- [9] D. Cheng, F. Yang, S. Xiang, J. Liu, Financial Time Series Forecasting with Multi-Modality Graph Neural Network, *Pattern Recognit.* 121 (2022), 108218. <https://doi.org/10.1016/j.patcog.2021.108218>.
- [10] F. Feng, X. He, X. Wang, C. Luo, Y. Liu, et al., Temporal Relational Ranking for Stock Prediction, *ACM Trans. Inf. Syst.* 37 (2019), 1–30. <https://doi.org/10.1145/3309547>.
- [11] C. Xu, H. Huang, X. Ying, J. Gao, Z. Li, et al., HGNN: Hierarchical Graph Neural Network for Predicting the Classification of Price-Limit-Hitting Stocks, *Inf. Sci.* 607 (2022), 783–798. <https://doi.org/10.1016/j.ins.2022.06.010>.
- [12] S. Saha, J. Gao, R. Gerlach, A Survey of the Application of Graph-Based Approaches in Stock Market Analysis and Prediction, *Int. J. Data Sci. Anal.* 14 (2022), 1–15. <https://doi.org/10.1007/s41060-021-00306-9>.
- [13] Y. Wang, Y. Qu, Z. Chen, Review of Graph Construction and Graph Learning in Stock Price Prediction, *Procedia Comput. Sci.* 214 (2022), 771–778. <https://doi.org/10.1016/j.procs.2022.11.240>.
- [14] R. Liu, H. Liu, H. Huang, B. Song, Q. Wu, Multimodal Multiscale Dynamic Graph Convolution Networks for Stock Price Prediction, *Pattern Recognit.* 149 (2024), 110211. <https://doi.org/10.1016/j.patcog.2023.110211>.
- [15] H. Qian, H. Zhou, Q. Zhao, H. Chen, H. Yao, et al., MDGNN: Multi-Relational Dynamic Graph Neural Network for Comprehensive and Dynamic Stock Investment Prediction, *Proc. AAAI Conf. Artif. Intell.* 38 (2024), 14642–14650. <https://doi.org/10.1609/aaai.v38i13.29381>.
- [16] D. Daiya, M. Yadav, H.S. Rao, Diffstock: Probabilistic Relational Stock Market Predictions Using Diffusion Models, in: *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2024, pp. 7335–7339. <https://doi.org/10.1109/ICASSP48485.2024.10446690>.

- [17] T.T. Huynh, B.T. Khoa, Financial Bubble Detection Using GSADF and LSTM-RNN Model: Evidence from Emerging Markets, *Int. J. Anal. Appl.* 23 (2025), 150. <https://doi.org/10.28924/2291-8639-23-2025-150>.
- [18] B.T. Khoa, T.T. Huynh, Long Short-Term Memory Recurrent Neural Network for Predicting the Return of Rate Underframe the Fama-French 5 Factor, *Discret. Dyn. Nat. Soc.* 2022 (2022), 3936122. <https://doi.org/10.1155/2022/3936122>.
- [19] N.T.K. An, Modeling and Forecasting Excess Market Returns Using a Two-Factor CIR Model, *Int. J. Anal. Appl.* 24 (2026), 180. <https://doi.org/10.28924/2291-8639-24-2026-180>.
- [20] T.T. Huynh, N.T.K. An, Dataset on Composite Efficiency Indices for ASEAN Foreign Exchange Markets, *Data Brief* 67 (2026), 112913. <https://doi.org/10.1016/j.dib.2026.112913>.