7. CNN-BASED STOCK PRICE FORECASTING BY STOCK CHART IMAGES

Jeongseok Bang[1]
Doojin Ryu[2]

Abstract

We use the recent development in deep learning technology to forecast stock prices. Focusing on image-type big data, we predict future stock prices using a convolutional neural network (CNN) model trained by visual representations of stock price data and technical indicators. We find that including technical indicators partially increases accuracy. The model with an input range of five days is the most accurate but is likely to be not appropriately learned, considering the recall, precision, and test datasets. On the contrary, training the model using past 20-day images along with technical indicators results in the greatest difference between the precision and label means of the test dataset.

Keywords: *Convolution neural networks; Stock chart image; Stock price forecasting; Technical indicators*

JEL Classification: *G11, G12, G17*

1. Introduction

One of the most important questions in financial research is whether the stock market can be systematically predicted. According to Fama's efficient market hypothesis, the stock market efficiently reflects the information in prices. In other words, the more efficient the market is, the faster it reaches equilibrium owing to instant arbitrage, and stock prices follow a random process. If the semi-strong efficient market hypothesis is correct, it is impossible to achieve systemic arbitrage using information publicly obtained from the market. However, following the work of Fama and French, evidence of anomalies in the stock market is increasing. Jegadeesh and Titman (1993) propose the concept of stock price momentum, suggesting that long-term price trends can generate excess returns that asset pricing models cannot explain.[3] This idea has led to the establishment of effective arbitrage strategies ranging from days to weeks (De Groot, Huij, and Zhou, 2012; Novy-Marx and Velikov, 2016).

Consequently, numerous studies have investigated whether the stock market dynamics can be explained or predicted based on cross-sectional analyses (Bang, Ryu, and Webb, 2023; Chen,

---

[1] *Department of Economics, Sungkyunkwan University, Seoul, Korea.*

[2] ***Corresponding author*** *(Doojin Ryu). Email: doojin.ryu@gmail.com*

Acknowledgment: *This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT; Ministry of Science and ICT) [No. 2022R1A2C1010596].*

[3] *In financial economics academic communities, their approaches are widely studied research topics in terms of the effectiveness of various cross-sectional and time-series momentum strategies (Lee, Batten, Ham, and Ryu, 2023; Ham, Cho, Kim, and Ryu, 2019).*

Han, Ryu, and Tang, 2022; Chun, Cho, and Ryu, 2023; George, Hwang, and Li, 2018; Ham, Ryu, and Webb, 2022; Kim and Ryu, 2022; Lee, Cho, Ryu, and Seok, 2023; Seok, Cho, Lee, and Ryu, 2023; Yu and Ryu, 2019). Other studies explore stock market forecasting by time series analysis, such as autoregressive integrated moving average, vector autoregressive, heterogeneous autoregressive, and autoregressive conditionally heteroskedastic models (Chun, Cho, and Ryu, 2019, 2020; Han, Kutan, and Ryu, 2015; Kim and Ryu, 2020; Park, Ryu, and Song, 2017; Song, Ryu, and Webb, 2016, 2018; Song, Park, and Ryu, 2018). However, due to the inherent complexity of stock markets, recent research has increasingly focused on non-linear approaches, particularly deep learning, to predict asset prices (Chen *et al.*, 2018; Nakano, Takahashi, and Takahashi, 2018; Shen and Shafiq, 2020). Recent financial studies practically and actually apply the new machine learning technologies to predict and explain financial market dynamics (Chen, Pelger, and Zhu, 2023; Gu, Kelly, and Xiu, 2020; Park and Ryu, 2021a, 2021b; Kim, Cho, and Ryu, 2021a, 2021b, 2022, 2023). Bustos and Pomares-Quimbaya (2020) use support vector machines, tree-ensemble models, and multiple perceptron models for stock market predictions. Furthermore, as the recurrent neural network models evolve, researchers develop stock market prediction models with long short-term memory and gated recurrent units using stock prices, technical indicators, and macroeconomic data as inputs (Jiang *et al.*, 2020). Meanwhile, as many investors predict future stock prices through stock price charts, studies also attempt to predict short-term price changes using chart images instead of numerical data (Chen *et al.*, 2021). If a significant number of investors recognize stock price movements as images and trade stocks, stock price forecasting by stock price images may be an effective method (Jiang, Kelly, and Xiu, 2023). If meaningful prediction patterns originated from price movements and numerous technical indicators exist, then they can be extracted using an image-based deep learning model.

In a related study, Chen *et al.* (2021) propose a graph convolutional feature-based convolutional neural network (GC-CNN) model that considers both individual stocks and the correlations between stocks. They find that GC-CNN-based strategies perform better than other models do. They analyze stock price charts and abstract images of technical indicators using a Daul-convolutional neural network model. Jiang, Kelly, and Xiu (2023) present a strategy to train CNNs with stock chart images and construct long-short portfolios based on the probability that stock prices will rise. They find that this strategy has a relatively high Sharpe ratio, outperforming other short-term momentum strategies, and is robust to transaction costs.

This study expands on the recent study of Jiang, Kelly, and Xiu (2023) to investigate potential stock price patterns. The patterns can be captured by a CNN model which is learned by stock price images. Thus, we consider well-known technical indicators, such as the Bollinger band and the relative strength index (RSI), in addition to the moving averages, which are expressed as images to examine whether adding technical indicators improves the model's predictability.[4] We train our CNN model using the opening, high, low, and closing prices; trading volumes; moving averages; Bollinger bands; and RSI in bar chart images with input windows of 5, 20, and 60 days. We find that forecasting the next five days using the prior five days performs better than the other input windows do. In addition, the trading volumes, moving averages, and Bollinger bands tend to provide more accurate predictions. However, when we analyze the recall, precision, and label mean of the test dataset, the high accuracy of five days is only due to the label mean of the test data, suggesting that learning the model with a 5-day input window may be spurious. Instead, using 20- and 60-day windows leads to more meaningful learning despite their relatively low accuracy.

[4] *These technical indicators are also widely used to construct sentiment indices in financial markets, which predict stock market dynamics (Kim and Ryu, 2021a, 2021b; Kim, Ryu, and Yang, 2021; Kim, Ryu, and Yu, 2021, 2022; Seok, Cho, and Ryu, 2021, 2022; Ryu, Ryu, and Yang, 2020, 2023).*

2. Data and Methodology

2.1. Stock Price Image Transformation

We collect stock price data for the largest 100 Nasdaq-listed companies, selected according to market capitalization as of June 2023, using the Yahoo Finance API. The data cover the period from January 1, 2000, to December 31, 2022. We convert the numeric data, including opening, high, low, and closing prices and trading volumes, into bar chart images. To this end, following Jiang, Kelly, and Xiu (2023), we convert numerical price data into a matrix between [0, 255]. We use 5-, 20-, and 60-day input windows and 5-day output windows. Through this process, we derive the images shown in Figure 1. Each image has a size of (256, input window size × 3). Of the three pixels assigned to each date, the first pixel represents the opening price, the second pixel shows a line connecting the low and high prices, and the third pixel represents the closing price. The price information is simplified, as it is mapped to a single integer out of the 256 possible values. The highest and lowest values in the given price series correspond to the top and bottom of the image, respectively.

Figure 1. Input images

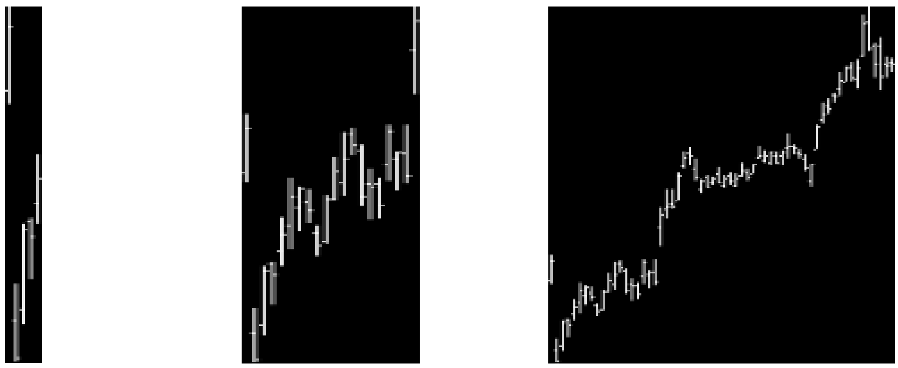

(a) 5-day window (b) 20-day window (c) 60-day window

The four versions of the images used to train the CNN model are shown in Figure 2. Version 1 represents the opening, high, low, and closing prices in bar chart formats. Version 2 depicts the prices and moving average lines of the closing prices for each input window in the top 75% of the image while the bottom 25% represents the trading volume data. Version 3 includes the upper and lower Bollinger bands (20-day windows) based on closing prices. The area between the upper and lower bands is shaded at 128 which is the midpoint between zero and 255. Version 4 follows the same approach as version 2, but each column is multiplied by the RSI (calculated on a 14-day basis), and the products represent brightness. Versions 3 and 4 allow us to incorporate additional technical indicators, Bollinger bands, and the RSI, into the images.

We construct the output data as follows. The output equals one if the stock price increases over the next five days and zero otherwise. The training data span the period from January 1, 2000, to December 31, 2018, and the test data cover the period from January 1, 2019, to December 31, 2022.

Figure 2. Input image data using technical indicators

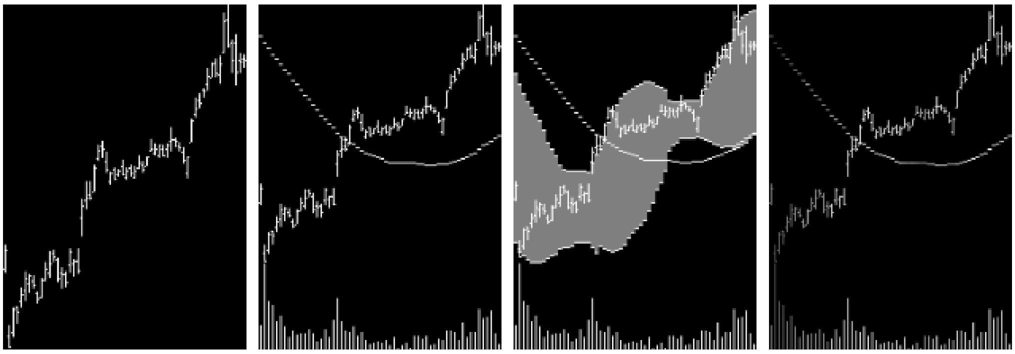

(a) Version 1 (b) Version 2 (c) Version 3 (d) Version 4

2.2. CNN Architecture

In this study, we use a CNN model due to the data being in image form. The structure of the CNN used in this study consists of two layers composed of convolution, activation functions, and pooling for each, as illustrated in Figure 3. In the first layer, we use the convolutional layer, a set of eight filters with a 3 × 3 shape, to train the model on the stock price chart images. This process captures the spatial patterns in the images. To preserve the original image shapes, we apply a padding of one. We also use a leaky ReLU with a gradient of 0.01 for the activation function. A max-pooling layer with a 2 × 2 filter and a stride of two is adapted to reduce the height and width of the image by half. We use the Adam optimizer with a batch size of 50, 10 training epochs, and a learning rate of 0.001.

Figure 3. CNN architecture

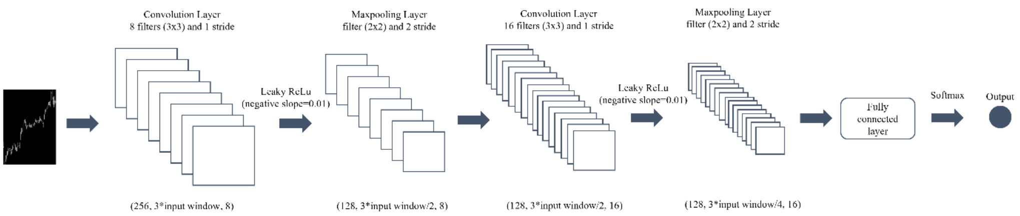

3. Results

Table 1 represents the number of image samples from the training and test data used in the analysis. Because the entire stock price time series for each firm is split by the given input window, the number of observations is largest when the input window is set to five days and is smallest when the input window is 60 days.

Table 1. The number of image samples in the training and testing data

	Version 1	Version 2	Version 3	Version 4
5 days	(83094, 19246)	(83830, 19444)	(83545, 19435)	(83756, 19444)
20 days	(20934, 4861)	(20840, 4858)	(20840, 4858)	(20840, 4858)
60 days	(6948, 1556)	(6853, 1553)	(6853, 1553)	(6853, 1553)

Table 2 presents the accuracy, precision, recall, and F1 scores of the trained CNN model using the input windows and the four image versions. The asterisks indicate the highest score for each metric. Overall, the 5-day input window provides the highest accuracy, and version 3, which includes Bollinger bands, is the most accurate version for this window. For both the 5-day and 60-day input windows, versions 2-4, which include additional technical indicators, are more accurate than version 1. In addition, for the 60-day input window, version 2 is the most accurate.

Table 2. Results of CNN models

		Accuracy	Precision	Recall	F1 Score
5 days	Version 1	0.5142	0.5482	0.9199*	0.6870*
	Version 2	0.5377	0.5468	0.8476	0.6648
	Version 3	0.5417*	0.5483	0.6646	0.6008
	Version 4	0.5304	0.5485*	0.8145	0.6555
20 days	Version 1	0.5069	0.4958	0.5932	0.5402
	Version 2	0.5010	0.5053	0.5510	0.5272
	Version 3	0.5037	0.5077	0.4669	0.4864
	Version 4	0.5023	0.4893	0.5902	0.5350
60 days	Version 1	0.4968	0.4932	0.5756	0.5312
	Version 2	0.5119	0.5016	0.6071	0.5493
	Version 3	0.5061	0.4833	0.5900	0.5314
	Version 4	0.5029	0.4977	0.5756	0.5338

Version 4 with a 5-day window achieves the highest value of precision that indicates how well the predicted classes are consistent with the real classes. For the 20-day input window, versions 2-3 with additional technical indicators provide more precision than version 1 does. For the 60-day input window, version 2, which includes moving averages and trading volumes, exhibits the highest precision. Recall measures how accurately the predicted classes match the actual classes for a certain class. The recall values for the 5-day window are close to one. However, for the 20-day and 60-day input windows, the recall values are around 0.4-0.5. Finally, the F1 score, which is the harmonic mean of precision and recall, tends to be higher for the 5-day input window.

Considering the overall accuracy, one may conclude that the CNN model performs best when the input window is five days. However, the extremely high recall values in the binary classification indicate that the model may learn to select only one class instead of well-fitted training. To validate the accuracy of the classes predicted using the trained CNN model, we compare them with the true ratio of label 1 in the test dataset. Table 3 presents the precision, the true ratio, and the differences between them. Versions 1-4 have relatively high precision for the 5-day input window, with relatively small differences ranging from 0.0007 to 0.0024. This result supports that the high accuracy observed in the model trained with 5-day images results from the statistical characteristics of the test data rather than proper model training. However, the differences in precision and true ratio for the 20-day and 60-day windows are relatively high, reaching a maximum of 0.0135. This result suggests that model training is not entirely meaningless despite its low accuracy. Version 3 with a 20-day image window, which has an accuracy of 0.5037, has the largest difference. Thus, despite being less accurate than the model using a 5-day input window, the models with longer image windows are likely to have been trained more meaningfully.

Table 3. Differences between the precision and the true ratio

		Precision	True ratio	Difference
5 days	Version 1	0.5482	0.5466	0.0016
	Version 2	0.5468	0.5461	0.0007
	Version 3	0.5483	0.5461	0.0022
	Version 4	0.5485	0.5461	0.0024
20 days	Version 1	0.4958	0.4956	0.0002
	Version 2	0.5053	0.4942	0.0111
	Version 3	0.5077	0.4942	0.0135
	Version 4	0.4893	0.4942	-0.0049
60 days	Version 1	0.4932	0.4891	0.0041
	Version 2	0.5016	0.4900	0.0116
	Version 3	0.4833	0.4900	0.0067
	Version 4	0.4977	0.4900	0.0077

4. Conclusion

Previous studies propose stock price forecasting with machine learning and deep learning techniques to effectively predict the stock market. Some of these approaches attempt to process stock price data as images and train deep learning models. This study represents stock movements as images and analyzes whether image-based deep-learning models can forecast future stock prices. We convert the stock price and technical indicator data into images for input windows of 5, 20, and 60 days and train CNN models using these images. As such, we try to uncover hidden patterns in stock price movements that are difficult to capture using traditional methodologies. The results show that the highest accuracy is achieved when the input window is 5 days. Considering the precision and the true ratio of testing data, however, learning over a shorter window may not extract potential patterns in stock price flows. Instead, we observe that 20- and 60-day windows lead to more meaningful learning. Including technical indicators improves accuracy partially.

References

Bang, J., Ryu, D., and Webb, R.I., 2023. ESG controversy as a potential asset-pricing factor. *Finance Research Letters*, 58(A), 104315. https://doi.org/10.1016/j.frl.2023.104315.

Bustos, O., and Pomares-Quimbaya, A., 2020. Stock market movement forecast: A systematic review. *Expert Systems with Applications*, 156, 113464. https://doi.org/10.1016/j.eswa.2020.113464.

Chen, L., Pelger, M., and Zhu, J., 2023. Deep learning in asset pricing. *Management Science*, *Forthcoming*. https://doi.org/10.1287/mnsc.2023.4695.

Chen, L., Qiao, Z., Wang, M., Wang, C., Du, R., and Stanley, H.E., 2018. Which artificial intelligence algorithm better predicts the Chinese stock market?. *IEEE Access*, 6, pp.48625-48633. https://doi.org/10.1109/ACCESS.2018.2859809.

Chen, J., Han, Q., Ryu, D., and Tang, J., 2022. Does the world smile together? A network analysis of global index option implied volatilities. *Journal of International Financial Markets, Institutions and Money*, 77, 101497. https://doi.org/10.1016/j.intfin.2021.101497.

Chen, W., Jiang, M., Zhang, W. G., and Chen, Z., 2021. A novel graph convolutional feature based convolutional neural network for stock trend prediction. *Information Sciences*, 556, pp.67-94. https://doi.org/10.1016/j.ins.2020.12.068.

Chun, D., Cho, H., and Ryu, D., 2019. Forecasting the KOSPI200 spot volatility using various volatility measures. *Physica A: Statistical Mechanics and its Applications*, 514, pp.156-166. https://doi.org/10.1016/j.physa.2018.09.027.

Chun, D., Cho, H., and Ryu, D., 2020. Economic indicators and stock market volatility in an emerging economy. *Economic Systems*, 44(2), 100788. ging economy. Economic Systems 44 (2), 100788. https://doi.org/10.1016/j.ecosys.2020.100788.

Chun, D., Cho, H., and Ryu, D., 2023. Discovering the drivers of stock market volatility in a data-rich world. *Journal of International Financial Markets, Institutions and Money*, 82, 101684. https://doi.org/10.1016/j.intfin.2022.101684.

De Groot, W., Huij, J., and Zhou, W., 2012. Another look at trading costs and short-term reversal profits. *Journal of Banking and Finance*, 36(2), pp.371-382. https://doi.org/10.1016/j.jbankfin.2011.07.015.

George, T. J., Hwang, C. Y., and Li, Y., 2018. The 52-week high, q-theory, and the cross section of stock returns. Journal of Financial Economics, 128(1), pp.148-163. https://doi.org/10.1016/j.jfineco.2018.01.005.

Gu, S., Kelly, B., and Xiu, D., 2020. Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), pp.2223-2273. https://doi.org/10.1093/rfs/hhaa009.

Ham, H., Cho, H., Kim, H., and Ryu, D., 2019. Time-series momentum in China's commodity futures market. *Journal of Futures Markets*, 39(12), pp.1515-1528. https://doi.org/10.1002/fut.22053.

Ham, H., Ryu, D., and Webb, R. I., 2022. The effects of overnight events on daytime trading sessions. *International Review of Financial Analysis*, 83, 102228. https://doi.org/10.1016/j.irfa.2022.102228.

Han, H., Kutan, A.M., and Ryu, D., 2015. Effects of the US stock market return and volatility on the VKOSPI. *Economics: The Open-Access, Open-Assessment Journal*, 9(1), 20150035. https://doi.org/10.5018/economics-ejournal.ja.2015-35.

Jegadeesh, N., and Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1), pp.65-91. https://doi.org/10.1111/j.1540-6261.1993.tb04702.x.

Jiang, J., Kelly, B.T., and Xiu, D., 2023. (Re-) Imag (in) ing price trends. *Journal of Finance*, *Forthcoming*. https://doi.org/10.1111/jofi.13268.

Jiang, M., Liu, J., Zhang, L., and Liu, C., 2020. An improved Stacking framework for stock index prediction by leveraging tree-based ensemble models and deep learning algorithms. *Physica A: Statistical Mechanics and its Applications*, 541, 122272. https://doi.org/10.1016/j.physa.2019.122272.

Kim, H., Cho, H., and Ryu, D., 2021a. Predicting corporate defaults using machine learning with geometric-lag variables. *Investment Analysts Journal*, 50(3), pp.161-175. https://doi.org/10.1080/10293523.2021.1941554.

Kim, H., Cho, H., and Ryu, D., 2021b. Forecasting consumer credit recovery failure: Classification approaches. *Journal of Credit Risk*, 17(3), pp.117-140. http://doi.org/10.21314/JCR.2021.007

Kim, H., Cho, H., and Ryu, D., 2022. Corporate bankruptcy prediction using machine learning methodologies with a focus on sequential data. *Computational Economics*, 59, pp.1231-1249. https://doi.org/10.1007/s10614-021-10126-5.

Kim, H., Cho, H., and Ryu, D., 2023. Measuring corporate failure risk: Does long short-term memory perform better in all markets?. *Investment Analysts Journal*, 52(1), pp.40-52. https://doi.org/10.1080/10293523.2022.2155353.

Kim, K., and Ryu, D., 2020. Predictive ability of investor sentiment for the stock market. *Romanian Journal of Economic Forecasting*, 23(4), pp.33-46.

Kim, K., and Ryu, D., 2021a. Does sentiment determine investor trading behaviour?. *Applied Economics Letters*, 28(10), pp.811-816. https://doi.org/10.1080/13504851.2020.1782331.

Kim, K., and Ryu, D., 2021b. Term structure of sentiment effect on investor trading behavior. *Finance Research Letters*, 43, 102005. https://doi.org/10.1016/j.frl.2021.102005.

Kim, K., Ryu, D., and Yang, H., 2021. Information uncertainty, investor sentiment, and analyst reports. *International Review of Financial Analysis*, 77, 101835. https://doi.org/10.1016/j.irfa.2021.101835

Kim, K., Ryu, D., and Yu, J., 2021. Do sentiment trades explain investor overconfidence around analyst recommendation revisions?. *Research in International Business and Finance*, 56, 101376. https://doi.org/10.1016/j.ribaf.2020.101376.

Kim, K., Ryu, D., and Yu, J., 2022. Is a sentiment-based trading strategy profitable?. *Investment Analysts Journal*, 51(2), pp.94-107. https://doi.org/10.1080/10293523.2022.2076373.

Kim, Y., and Ryu, D., 2022. Firm-specific or market-wide information: How does analyst coverage influence stock price synchronicity?. *Borsa Istanbul Review*, 22(6), pp.1069-1078. https://doi.org/10.1016/j.bir.2022.07.011.

Lee, J., Batten, J.A., Ham, H., and Ryu, D., 2023. Does portfolio momentum beat analyst advice?. *Abacus*, *Forthcoming*. https://doi.org/10.1111/abac.12300.

Lee, J.E., Cho, H., Ryu, D., and Seok, S., 2023. Does performance-chasing behavior matter? International evidence. *Journal of Multinational Financial Management*, 68, 100799. https://doi.org/10.1016/j.mulfin.2023.100799.

Nakano, M., Takahashi, A., and Takahashi, S., 2018. Bitcoin technical trading with artificial neural network. *Physica A: Statistical Mechanics and its Applications*, 510, pp.587-609. https://doi.org/10.1016/j.physa.2018.07.017.

Novy-Marx, R., and Velikov, M., 2016. A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1), pp.104-147. https://doi.org/10.1093/rfs/hhv063.

Park, D., and Ryu, D., 2021a. A machine learning-based early warning system for the housing and stock markets. *IEEE Access*, 9, pp.85566-85572. https://doi.org/10.1109/ACCESS.2021.3077962

Park, D., and Ryu, D., 2021b. Forecasting stock market dynamics using bidirectional long short-term memory. *Romanian Journal of Economic Forecasting*, 24(2), pp.22-34.

Park S.Y., Ryu, D., and Song, J., 2017. The dynamic conditional relationship between stock market returns and implied volatility. *Physica A: Statistical Mechanics and its Applications*, 482, pp.638-648. https://doi.org/10.1016/j.physa.2017.04.023.

Ryu, D., Ryu, D., and Yang, H., 2020. Investor sentiment, market competition, and financial crisis: Evidence from the Korean stock market. *Emerging Markets Finance and Trade*, 56(8), pp.1804-1816. https://doi.org/10.1080/1540496X.2019.1675152.

Ryu, D., Ryu, D., and Yang, H., 2023. Whose sentiment explains implied volatility change and smile?. *Finance Research Letters*, 55(A), 103838. https://doi.org/10.1016/j.frl.2023.103838.

Shen, J., and Shafiq, M.O., 2020. Short-term stock market price trend prediction using a comprehensive deep learning system. *Journal of Big Data*, 7(1), pp.1-33. https://doi.org/10.1186/s40537-020-00333-6.

Seok, S., Cho, H., Lee, J.E., and Ryu, D., 2023. Indirect effects of flow-performance sensitivity on fund performance. *Borsa Istanbul Review*. https://doi.org/10.1016/j.bir.2023.08.002.

Seok, S.I., Cho, H., and Ryu, D., 2021. Stock market's responses to intraday investor sentiment. *North American Journal of Economics and Finance*, 58, 101516. https://doi.org/10.1016/j.najef.2021.101516.

Seok, S.I., Cho, H., and Ryu, D., 2022. Scheduled macroeconomic news announcements and intraday market sentiment. *North American Journal of Economics and Finance*, 62, 101739. https://doi.org/10.1016/j.najef.2022.101739.

Song, W., Park, S.Y., and Ryu, D., 2018. Dynamic conditional relationships between developed and emerging markets. *Physica A: Statistical Mechanics and its Applications*, 507, pp.534-543. https://doi.org/10.1016/j.physa.2018.05.007.

Song, W., Ryu, D., and Webb R.I., 2016. Overseas market shocks and VKOSPI dynamics: A Markov-switching approach. *Finance Research Letters*, 16, pp.275-282. https://doi.org/10.1016/j.frl.2015.12.007.

Song, W., Ryu, D., and Webb, R.I., 2018. Volatility dynamics under an endogenous Markov-switching framework: A cross-market approach. *Quantitative Finance*, 18(9), pp.1559-1571. https://doi.org/10.1080/14697688.2018.1444551.

Yu, J., and Ryu, D., 2019. Predicting banks' subordinated bond issuances. *Romanian Journal of Economic Forecasting*, 22(4), pp.87-99.