

Detecting Anomalies in Foreign Exchange Markets: A Comparison between Structural and Generative Approaches

Andreea-Madalina BOZAGIU

Bucharest University of Economic Studies, Bucharest, Romania

andreea.bozagiu@fin.ase.ro

Abstract. *This paper investigates and compares two distinct approaches to anomaly detection in foreign exchange markets: a structural approach based on correlation networks (Machine Learning approach) and a generative reconstruction-based approach (Generative AI approach), using a twenty-year sample of returns for the EUR–RON, EUR–HUF, EUR–PLN, and EUR–USD exchange rate pairs. The analysis examines whether the two approaches identify the same anomalous periods and what each approach captures. The results show that graph-based methods mainly detect anomalies linked to changes in the structure of currency relationships, even when price movements are not extreme. In contrast, generative models are more sensitive to local shocks in returns but less effective at capturing changes in cross-currency relationships. The paper offers a direct comparison of the two approaches and shows that relational information plays an important role in understanding systemic behavior in foreign exchange markets.*

Keywords: Foreign exchange markets, Machine Learning, Anomaly detection, Network analysis, Generative AI.

Introduction

The main reason for choosing this topic is the increasing complexity of foreign exchange markets and the need to identify market anomalies in a reliable way. Exchange rate dynamics are influenced by a various range of macroeconomic, financial, and geopolitical factors which can lead to abrupt changes in market behavior, such as increased volatility or stronger co-movements between currencies. Detecting such anomalies is essential for risk management, portfolio allocation, and financial stability analysis, especially in markets characterized by strong cross-currency interactions.

Traditionally, anomaly detection in foreign exchange markets has relied on return-based indicators, volatility measures, or threshold rules applied to individual exchange rates. While these approaches are effective in identifying extreme price movements, they often fail to capture abnormal behavior arising from changes in the relationships between currencies. More recently, Machine Learning techniques have been increasingly adopted for anomaly detection, including unsupervised methods and generative models, which are designed to identify unusual patterns in high-dimensional data without requiring labeled observations. However, many of these methods focus on return dynamics and provide limited insight into the structural organization of the market.

In parallel, network-based approaches have gained attention in financial economics to model interdependencies between assets. By representing financial markets as networks of correlated assets, these methods allow for the analysis of market structure, connectivity, and synchronization. Despite their growing use in systemic risk and market integration studies, network-based techniques remain relatively underexplored in the context of anomaly detection, particularly in comparison with modern generative models.

The present study aims to compare two different approaches to anomaly detection in foreign exchange markets: a graph-based structural method and a generative model based on an

autoencoder. The analysis focuses on the EUR–RON, EUR–HUF, EUR–PLN, and EUR–USD exchange rate pairs and investigates whether anomalies identified through changes in correlation structures differ from those detected through return-based reconstruction errors.

Despite the growing literature on anomaly detection in financial markets, existing studies tend to focus either on return-based machine learning techniques or on network-based representations of market structure, with limited interaction between the two strands. Applications to foreign exchange markets typically adopt a single methodological perspective, making it difficult to assess how different approaches capture distinct forms of abnormal market behavior. This paper addresses this gap by providing a direct and systematic comparison between a graph-based structural anomaly detection framework and a generative reconstruction-based model applied to the same set of exchange rates and observation periods. The results offer practical insights for researchers and practitioners seeking to improve the detection and interpretation of abnormal market conditions.

The contribution of this study is threefold. First, it demonstrates that structural, network-based methods can identify anomalies associated with changes in cross-currency dependency patterns, even in the absence of extreme price movements. Second, it shows that generative autoencoder models primarily detect localized return shocks and may overlook broader structural regime shifts. Third, by jointly analyzing and contrasting the outputs of the two approaches, the paper highlights the complementary nature of structural and generative anomaly detection methods and provides insights into their relative interpretability and economic relevance in foreign exchange markets.

Literature review

Anomaly detection in financial markets has attracted considerable attention due to its relevance for risk management, market surveillance, and financial stability analysis. Early contributions in this area primarily focused on identifying extreme observations in individual asset returns using statistical techniques based on distributional assumptions. Seminal studies by Barnett and Lewis (1994) and Tsay (2005) rely on univariate outlier detection methods to identify abnormal price movements. While effective in detecting extreme events, these approaches are limited by their inability to capture cross-asset dependencies, which play a central role in foreign exchange markets.

Further research extends anomaly detection to multivariate settings by incorporating information from multiple assets simultaneously. Methods based on principal component analysis and distance-based techniques are discussed by Jolliffe (2002) and Aggarwal (2013), who show that multivariate approaches improve detection accuracy in high-dimensional financial data. Machine learning techniques, including clustering and classification models, are further surveyed by Chandola, Banerjee, and Kumar (2009). Although these approaches account for joint behavior across assets, dependencies are typically treated implicitly, and the evolving structure of relationships between financial instruments is not explicitly modeled.

A separate strand of the literature applies network theory to financial markets by modeling assets as nodes and their interactions as edges derived from correlation measures. Mantegna (1999) provides one of the earliest applications of correlation-based networks in finance, demonstrating that financial markets exhibit non-random and hierarchical structures. Later studies by Kwapien and Drozd (2012) and Diebold and Yilmaz (2014) show that network topology changes significantly during periods of financial stress, with increased connectivity and synchronization among assets. In the context of foreign exchange markets, McDonald et al. (2005) and Mizuno et

al. (2006) document that currency networks become more densely connected during turbulent periods, suggesting that structural changes in dependency patterns may serve as indicators of systemic risk.

In parallel, generative models have gained prominence as tools for anomaly detection in financial time series. Autoencoders and related architectures are discussed by Sakurada and Yairi (2014) and An and Cho (2015), who show that reconstruction-based methods are effective in identifying rare or unusual patterns in multivariate data. More recent contributions, such as Malhotra et al. (2016), apply recurrent autoencoders to financial time series and demonstrate strong performance in detecting localized shocks. However, these models primarily focus on learning typical return dynamics and do not explicitly incorporate relational information between assets, which limits their ability to detect anomalies driven by changes in cross-asset dependency structures.

Recent studies have also emphasized the importance of adopting comparative and hybrid perspectives in anomaly detection. Akoglu, Tong, and Koutra (2015) argue that different anomaly detection techniques capture fundamentally different notions of abnormality, ranging from point anomalies to contextual and structural anomalies. In a financial setting, Buczak and Guven (2016) similarly highlight that no single method is sufficient to fully characterize abnormal market behavior, as different models respond to distinct features of the data. These insights motivate comparative analyses that evaluate how alternative methodological frameworks differ in terms of sensitivity to market conditions, interpretability, and economic relevance.

Overall, the literature points to two complementary approaches to anomaly detection in financial markets: structural methods that emphasize evolving relationships between assets and generative methods that focus on deviations from learned patterns. Despite extensive research in both areas, direct comparisons between network-based structural approaches and generative reconstruction-based models in foreign exchange markets remain limited. This paper contributes to the literature by addressing this gap and providing a systematic comparison of graph-based structural anomaly detection and generative autoencoder-based methods, offering new insights into the nature of anomalies and systemic behavior in currency markets.

Methodology

This study employs two complementary approaches to anomaly detection in foreign exchange markets: a graph-based structural method and a generative method based on an autoencoder. Both approaches are applied to the same dataset and observation windows to ensure comparability and to highlight the different types of abnormal market behavior captured by each framework.

Data and Preprocessing

This study employs two distinct approaches for detecting anomalies in foreign exchange markets to compare their effectiveness: a graph-based structural anomaly detection method and a generative anomaly detection method based on autoencoders. The analysis focuses on daily exchange rate data for the EUR–RON, EUR–HUF, EUR–PLN, and EUR–USD currency pairs starting with 3rd of January 2005 to 3rd of December 2025. Both approaches are applied to the same dataset to ensure comparability of results. Prior to analysis, exchange rate series are transformed into logarithmic returns to reduce non-stationarity and ensure comparability across currency pairs. This transformation allows the focus to be placed on relative changes in exchange rates rather than on their absolute levels.

Graph-Based Structural Anomaly Detection

To capture time-varying dependencies between currency pairs, the foreign exchange market is represented as a sequence of correlation networks. For each observation period, a correlation matrix C_t is computed from the return series, where each element $\rho_{ij,t}$ represents the Pearson correlation coefficient between currency pairs i and j :

$$\rho_{ij,t} = \frac{\text{Cov}(r_{i,t}, r_{j,t})}{\sigma_{i,t} \sigma_{j,t}} \quad (1)$$

Although alternative dependence measures could be employed, Pearson correlations are chosen due to their transparency, interpretability, and widespread use in the financial network literature.

For interpretational purposes, correlations can be transformed into distances using the standard metric proposed by Mantegna (1999):

$$d_{ij,t} = \sqrt{2(1 - \rho_{ij,t})} \quad (2)$$

Based on the correlation matrix, an undirected weighted network $G_t = (V, E_t)$ is constructed, where nodes represent currency pairs and edges capture the strength of linear dependence between them. To focus on economically meaningful relationships and reduce noise, edges are retained only when the absolute value of the correlation coefficient exceeds a predefined threshold.

To characterize the topology of each network, several complementary metrics are computed. Network density is defined as:

$$D_t = \frac{2|E_t|}{N(N-1)} \quad (3)$$

where $|E_t|$ is the number of edges and N is the number of nodes. The average degree is given by:

$$\bar{k}_t = \frac{1}{N} \sum_{i=1}^N k_{i,t} \quad (4)$$

Local synchronization is captured through the clustering coefficient. For a given node i :

$$C_{i,t} = \frac{2e_{i,t}}{k_{i,t}(k_{i,t}-1)} \quad (5)$$

and the average clustering coefficient of the network is:

$$C_t = \frac{1}{N} \sum_{i=1}^N C_{i,t} \quad (6)$$

Systemic integration is further assessed using the relative size of the largest connected component:

$$\text{LCC}_t = \frac{|V_t^{\max}|}{N} \quad (7)$$

Structural anomalies are identified by analyzing the joint behavior of these network metrics. Let $\mathbf{m}_t$ denote the vector of network indicators at time t . An anomaly score is defined as the multivariate distance from the typical network configuration:

$$A_t = \sqrt{(\mathbf{m}_t - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{m}_t - \boldsymbol{\mu})} \quad (8)$$

Periods associated with unusually high values of A_t are classified as structural anomalies, reflecting significant changes in the dependency structure of the foreign exchange market rather than isolated price movements.

Generative Anomaly Detection Using Autoencoders

In parallel, a generative anomaly detection approach based on an autoencoder neural network is employed. The autoencoder is trained on multivariate return vectors $\mathbf{r}_t$ and consists of an encoder that maps the input data into a lower-dimensional latent representation and a decoder that reconstructs the original input:

$$\mathbf{z}_t = f(\mathbf{r}_t), \hat{\mathbf{r}}_t = g(\mathbf{z}_t) \quad (9)$$

The model is trained by minimizing the reconstruction error:

$$\mathcal{L}_t = \|\mathbf{r}_t - \hat{\mathbf{r}}_t\|^2 \quad (10)$$

Anomaly detection is based on the distribution of reconstruction errors. A threshold τ is defined as a high quantile of empirical loss distribution:

$$\tau = Q_{1-\alpha}(\mathcal{L}_t) \quad (11)$$

Observations are classified as anomalous when:

$$\text{Anomaly}_t = \begin{cases} 1, & \mathcal{L}_t > \tau \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

This reconstruction-based criterion allows the identification of localized shocks and atypical return configurations without explicitly modeling cross-currency dependencies.

Comparative Framework and Robustness Considerations

To assess differences between the two approaches, anomalous periods identified by the graph-based method and the autoencoder are systematically compared across the full sample. The analysis examines both overlapping anomalies and periods uniquely identified by each method. Structural anomalies are interpreted as indicators of changes in cross-currency dependency structures, while generative anomalies are associated primarily with unusual return behavior.

The use of two conceptually distinct anomaly detection frameworks provides a robustness check of results. Periods detected by both approaches are interpreted as episodes of heightened market stress affecting both return dynamics and dependency structures, while discrepancies between methods offer insights into the specific dimensions of abnormal market behavior captured by each model.

Results and discussions

The first step of the analyses is to observe the exchange rate dynamics, followed by an examination of structural and generative anomaly detection outcomes. This approach allows a clear separation between baseline market behavior and periods identified as anomalous by the proposed methods.

Figure 1 illustrates the evolution of the EUR–RON, EUR–HUF, EUR–PLN, and EUR–USD exchange rate pairs over the sample period, normalized to an index value of 100 at the beginning of the series. This normalization facilitates a direct comparison of relative movements across currency pairs, despite differences in their nominal levels.

Several distinct patterns emerge from the figure. The EUR–HUF exchange rate exhibits the strongest upward trend and the highest volatility, indicating sustained depreciation of the Hungarian forint against the euro over the sample period. In contrast, the EUR–PLN series shows comparatively moderate fluctuations, suggesting a more stable exchange rate dynamic. The EUR–RON exchange rate displays a gradual upward movement with relatively low short-term volatility,

reflecting a smoother adjustment path. The EUR–USD series follows a markedly different trajectory, characterized by pronounced cycles and periods of depreciation and appreciation, highlighting its distinct global role relative to regional European currencies.

Importantly, the figure suggests that while exchange rates follow heterogeneous long-term paths, they also have periods of joint movements, particularly during episodes of market stress. These observations motivate a deeper analysis of the relationships between currency pairs and provide a visual justification for modeling the foreign exchange market as a system of interconnected assets rather than as independent time series.




Figure 1. Exchange rates evolution

Source: Authors' owns research.

However, such a representation does not capture how currencies move relative to one another over time. Periods in which exchange rates become more closely aligned may reflect broader economic or financial stress, even if individual price movements remain moderate. Figure 2 addresses this limitation by focusing on changes in the relationship between exchange rates. The anomaly score identifies periods in which currencies exhibit unusually strong or altered co-movements, providing an economically meaningful perspective on market conditions that is not directly observable from exchange rate levels alone.




Figure 2. Anomaly score over time

Source: Authors' owns research.

Figure 2 presents the evolution of the anomaly score obtained from the graph-based structural anomaly detection method. Higher values of the anomaly score indicate periods in which the underlying correlation structure of the foreign exchange market deviates more strongly from its typical behavior.

The figure reveals that anomalous periods are not uniformly distributed over time but tend to occur in clusters. Several pronounced peaks can be observed during specific intervals, indicating episodes in which the network of currency correlations undergoes substantial structural changes. These peaks do not necessarily coincide with extreme movements in individual exchange rates, suggesting that the detected anomalies capture changes in market structure rather than isolated price shocks.

Notably, the anomaly score remains elevated during periods commonly associated with heightened market uncertainty, while relatively stable periods exhibit lower and less volatile scores. This pattern supports the interpretation that the graph-based approach is sensitive to shifts in cross-currency relationships and synchronization. The presence of frequent, short-lived spikes also indicates that structural changes can occur rapidly and may revert quickly, highlighting the dynamic nature of foreign exchange markets. So it is demonstrated that the proposed structural anomaly detection method provides a time-varying measure of market stress based on relational information, offering insights that complement traditional volatility- or return-based indicators.

To better understand the nature of these anomalies, it is necessary to examine the configuration of currency correlations in more detail. Figure 3 provides this insight by comparing the correlation networks observed during a representative normal period and an anomalous period. This comparison allows for a direct visualization of how economic linkages between currencies strengthen, weaken, or change sign when the market transitions from a stable regime to an anomalous one.




Figure 3. Normal vs Anomalous correlation networks

Source: Authors' owns research.

Figure 3 compares the correlation networks corresponding to a representative normal period and an anomalous period identified by the graph-based anomaly detection method. Edges are displayed only when the absolute value of the correlation exceeds the chosen threshold (0.35), allowing the visualization to focus on economically meaningful relationships between currency pairs.

During the normal period, the network exhibits a relatively sparse structure, with moderate correlations among the Central and Eastern European currency pairs and consistently negative correlations between the euro-based currencies and the United States dollar. This configuration reflects a typical market environment in which regional currencies share some common dynamics while maintaining a degree of independence, and the dollar plays a distinct global role. The EUR–RON exchange rate appears as an isolated node, reflecting weak correlations with other currency pairs that remain below the chosen threshold. This indicates a relatively independent behavior under stable market conditions. In contrast, during the anomalous period, EUR–RON becomes connected to the regional network, suggesting increased synchronization during market stress.

In contrast, the anomalous period is characterized by a denser and more synchronized network structure. Correlations among regional currencies increase markedly, particularly between the EUR–HUF and EUR–PLN pairs, indicating stronger co-movements and reduced diversification within the region. At the same time, the negative relationship between the euro-based currencies and the United States dollar remains pronounced, reinforcing the role of the dollar as an external anchor during periods of market stress. The emergence of stronger and more numerous connections highlights a structural reorganization of the foreign exchange market rather than isolated abnormal price movements.

Below in Figure 4, the correlation matrices complement the network representation by reporting the full correlations for the same normal and anomalous periods, allowing for a direct quantitative comparison of cross-currency relationships. The reported dates correspond to the end of the observation periods used to characterize market conditions.




Figure 4. Heatmap correlation Normal vs Anomalous

Source: Authors' owns research.

The values reported in each cell correspond to pairwise correlation coefficients, providing a detailed numerical view of the relationships illustrated in the network graphs.

During the normal period, correlations among the Central and Eastern European currency pairs remain moderate. For example, the correlation between EUR–HUF and EUR–PLN is approximately 0.55, while the correlations involving EUR–RON are weaker, indicating a relatively diversified regional structure. The EUR–USD pair exhibits consistently negative correlations with the regional currencies, reflecting its distinct behavior relative to the European market.

In contrast, the anomalous period is characterized by a marked increase in regional synchronization. The correlation between EUR–HUF and EUR–PLN rises to approximately 0.77, while correlations involving EUR–RON also increase noticeably. These changes indicate a stronger alignment of regional currencies and a reduction in diversification. At the same time, the negative correlations between EUR–USD and the euro-based currency pairs persist, although their magnitudes change, suggesting that the role of the United States dollar remains structurally different even during anomalous periods.

So, Figure 4 provides quantitative evidence that anomalous periods detected by the proposed method are associated with significant changes in cross-currency dependencies rather than isolated return shocks. The heatmaps confirm that the identified anomalies reflect structural reorganization in the foreign exchange market, reinforcing the findings obtained from the network-based analysis.

The results reported in Table 1 represent a summary of the anomaly detection outputs obtained for all observation periods. For each period, a set of network-level indicators, including density, average degree, clustering coefficient, degree centralization, and the size of the largest connected component, was computed together with an overall anomaly score. Based on these outputs, periods were classified as anomalous using both the graph-based approach and the generative model.

Table 1. Top 5 anomaly detection periods comparison with generative model

Window end date	Graph-based anomaly	Generative AI anomaly	Dominant signal
2020-03-18	Yes	Yes	Structural + volatility
2020-04-02	Yes	Yes	Structural persistence
2022-02-28	Yes	No	Structural reorganization
2022-03-10	Yes	No	Correlation regime shift
2023-03-15	No	Yes	Localized return shock

Source: Authors' owns research.

Given the large number of observation periods, only a limited number of representative cases are reported. Table 1 includes the highest-ranked anomalous periods as well as selected cases detected by only one of the two methods, to illustrate the different types of anomalies captured by each approach.

The comparison reported in Table 1 also indicates that the generative model is able to identify anomalous periods that largely overlap with those detected by the graph-based approach. This suggests that generative anomaly detection can capture abnormal market behavior with a performance comparable to traditional machine learning techniques, despite relying on a fundamentally different mechanism. By learning typical patterns in exchange rate returns, the generative model can signal deviations that coincide with periods of heightened market stress.

From a practical perspective, this result highlights the potential of generative models as early warning tools. Because such models operate directly on return dynamics and do not require explicit modeling of market structure, they can provide timely signals of emerging instability. While they may be less effective at capturing structural reorganization in currency relationships, generative models can serve as a complementary system for early detection of abnormal market conditions, including the onset of financial crises.

Conclusion

This paper applies anomaly detection methods to foreign exchange markets using two complementary approaches: a graph-based structural method and a generative model based on an autoencoder. Focusing on the EUR–RON, EUR–HUF, EUR–PLN, and EUR–USD exchange rate pairs, the study highlights the importance of cross-currency relationships in identifying abnormal market conditions. Instead of focusing only on extreme price movements, the proposed framework detects anomalies related to changes in how exchange rates move together.

The empirical results show that the graph-based approach is particularly effective at detecting anomalies associated with structural reorganization and increased synchronization among currencies. These periods are characterized by denser correlation networks and stronger co-movements, indicating a temporary loss of diversification within the market. Importantly, such structural anomalies are not always accompanied by extreme fluctuations in individual exchange rates, suggesting that traditional return-based indicators may fail to capture early signs of systemic stress.

In contrast, the generative model primarily identifies anomalies linked to localized shocks in return dynamics. By learning typical patterns of multivariate returns, the autoencoder effectively flags periods of unusual price behavior. While this approach does not explicitly model interdependencies between currency pairs and is therefore less sensitive to structural regime shifts, the results show that it is able to detect anomalous periods that substantially overlap with those identified by the graph-based method.

From a practical and economic perspective, this finding suggests that generative models can perform comparably to traditional machine learning techniques in identifying abnormal market conditions. Due to their ability to operate directly on return dynamics and provide timely signals, generative models may serve as effective early warning tools for emerging market stress or the onset of financial crises. Although they offer less interpretability regarding structural changes, their responsiveness makes them valuable complements to network-based approaches.

To sum up, the results indicate that combining structural and generative methods can improve the robustness and timeliness of anomaly detection in foreign exchange markets. Network-based models offer interpretability and insights into systemic reorganization, while generative models contribute early detection of abnormal return behavior. Future research may extend this framework to a broader set of currencies, alternative dependency measures, or hybrid models that integrate structural indicators with generative signals to enhance early detection of market instability and regime transitions.

This paper contributes to the literature on anomaly detection in financial markets in several ways. First, it provides a direct comparison between graph-based structural methods and generative anomaly detection models applied to foreign exchange data, highlighting how each approach captures different types of abnormal market behavior. Second, the study demonstrates that network-based indicators can reveal structural changes in cross-currency relationships that are not detectable through return-based methods alone. Third, the results show that generative models can identify anomalous periods with performance comparable to traditional machine learning techniques, supporting their potential use as complementary early warning tools. By integrating structural and generative perspectives, this paper extends existing approaches to anomaly detection and offers a framework that improves both interpretability and timeliness in identifying abnormal conditions in foreign exchange markets.

Acknowledgments: During the preparation of this work, the author utilized Claude and ChatGPT 4 to enhance readability and language clarity. Following the use of this tool, the author meticulously reviewed and revised the content as necessary and assumed full responsibility for the final content of the publication

References

- Aggarwal, C. C. (2013). Outlier analysis. New York, NY: Springer.
- Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3), 626–688.
<https://doi.org/10.1007/s10618-014-0365-y>
- An, J., & Cho, S. (2015). Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, 2, 1–18.
- Barnett, V., & Lewis, T. (1994). Outliers in statistical data (3rd ed.). Chichester, UK: Wiley.
- Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176.
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134.
<https://doi.org/10.1016/j.jeconom.2014.04.012>
- Jolliffe, I. T. (2002). Principal component analysis (2nd ed.). New York, NY: Springer.
- Kwapien, J., & Drozd, S. (2012). Physical approach to complex systems. *Physics Reports*, 515(34), 115–226.
- Malhotra, P., Vig, L., Shroff, G., & Agarwal, P. (2016). Long short term memory networks for anomaly detection in time series. *Proceedings of the 23rd European Symposium on Artificial Neural Networks*, 89–94.
- Mantegna, R. N. (1999). Hierarchical structure in financial markets. *The European Physical Journal B*, 11(1), 193–197.
- McDonald, M., Suleman, O., Williams, S., Howison, S., & Johnson, N. F. (2005). Detecting a currency's dominance or dependence using foreign exchange network trees. *Physical Review E*, 72(4), 046106.
- Mizuno, T., Takayasu, H., & Takayasu, M. (2006). Correlation networks among currencies. *Physica A: Statistical Mechanics and Its Applications*, 364, 336–342.
- Sakurada, M., & Yairi, T. (2014). Anomaly detection using autoencoders with nonlinear dimensionality reduction. *Proceedings of the MLSDA 2014 Workshop*, 4–11.
- Tsay, R. S. (2005). Analysis of financial time series (2nd ed.). Hoboken, NJ: Wiley.