MALAYSIAN JOURNAL OF SCIENCE
MJS
ISSN 1394 - 3065

5th International Conference on Applied & Industrial Mathematics and Statistics 2025 (ICoAIMS 2025)

Special Issue

Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026) https://mjs.um.edu.my

Using Time-Series Anomaly Detection Approach for Identifying Arbitrage Opportunities

Jose Abad[a], Darren Pabiloña[a], Stephen Ty[a], and Jeric Briones[a,b*]

Abstract: Since market anomalies are tedious and difficult to manually identify, the use of anomaly detection techniques to find statistical arbitrage opportunities is becoming increasingly popular. As such, this study compares the performance of different pairs trading strategies on the top 100 Philippine stocks listed on the Philippine Stock Exchange to assess the feasibility of using machine learning-based anomaly detection models. Specifically, time series anomaly detection through generative adversarial networks and auto-encoder with regression were employed to detect statistical arbitrage opportunities. The comparison across the strategies was limited to the top five cointegrated pairs (of order one) with the smallest sum of squared residuals. Results show that, even after accounting for transaction costs, AER outperformed the benchmarks (distance and cointegration approaches) by a wide margin, while TadGAN achieved a slightly better performance than the benchmarks. These findings suggest that anomaly detection models have the potential to yield higher profits than conventional approaches.

Keywords: anomaly detection, machine learning, pairs trading, Philippine stocks, statistical arbitrage

1. Introduction

As most transactions in financial markets carry some level of risk, strategies have been formulated to minimize said risk. One commonly used strategy involves taking advantage of arbitrage. Arbitrage refers to the practice of exploiting price differences in markets to buy and sell assets for a profit. For example, one type of arbitrage locks in a *riskless* profit by (simultaneously) entering into transactions in two or more markets. To illustrate, consider a stock that is listed in both the New York Stock Exchange (NYSE) and the London Stock Exchange (LSE). Suppose the stock price is $12 in NYSE and £10 in LSE, and that the prevailing exchange rate is $1.25 per £1. To take advantage of this arbitrage opportunity, a trader can buy shares (for example, 1,000 shares, costing $12,000) in NYSE, sell the shares in LSE (earning £10,000), and convert the proceeds of the sale from British pounds to U.S. dollars (equivalent to $12,500). This transaction would net $500 in profit for the trader.

Another type of arbitrage is statistical arbitrage, which uses price differences indicated by statistical testing. Due to the nature of this strategy, it carries minimal risk. An example of this arbitrage strategy is *pairs trading*. Pairs trading is a form of statistical arbitrage that hedges risk by trading securities that show a strong historical correlation of moving together. It relies on the assumption of mean reversion between the prices of such pairs (Carrasco Blázquez et al., 2018). For example, pairs trading in the stock market involves three steps: identifying a pair of historically highly correlated stocks, monitoring the price spread between the stocks, and trading according to the magnitude of their price spread. In literature, two established methods for performing these steps are the *distance* and the *cointegration* approaches. The distance approach conducts pairs trading by choosing pairs with the smallest sum of squared deviations between their cumulative return indices, which indicates that these prices "move together" (Gatev et al., 2006). The cointegration approach trades pairs that exhibit cointegration, which is the property that a linear combination of two $I(1)$ stock prices produces a stationary time series (Krauss, 2017). Here, $I(1)$ refers to a time series process whose order of integration is one.

While arbitrage strategies can be used to minimize or mitigate risks in transactions, their effectiveness relies on the ability of traders to detect opportunities to employ them. These opportunities are often present whenever market anomalies exist. Market anomalies refer to patterns in a market, such as the stock market, that contradict the efficient market hypothesis (Canady, 2022). Although they are difficult to identify manually, recent studies have shown that machine learning and artificial intelligence can be potentially used to detect them. Machine learning methods have exhibited viability in solving complex problems in other disciplines due to their ability to detect patterns even in complex data. This is especially evident with issues that might involve features that are tedious or difficult to manually solve or identify. For instance, in network intrusion detection systems, researchers have found that machine learning models and deep neural networks can significantly improve detection rates and reduce false positives compared to conventional approaches (Ullah et al., 2024). Of interest to this work are machine learning models used for anomaly detection. For example, deep learning models, such as the Transformer model, have been considered for time series anomaly detection

Authors information:
[a]Department of Mathematics, Ateneo de Manila University, Quezon City, PHILIPPINES
[b]Department of Finance and Accounting, Ateneo de Manila University, Quezon City, PHILIPPINES

*Corresponding Author: jbriones@ateneo.edu

Received: October 28, 2025
Accepted: May 6, 2026
Published: November 15, 2026

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)

(Nakashima & Yairi, 2023). However, some of these models, including the Transformer model, have not performed as expected. TadGAN, built on generative adversarial networks (GANs), and auto-encoder with regression (AER), however, have exhibited improvements in anomaly detection for time series data (Geiger et al., 2020; Wong et al., 2022). As such, the focus of this work is on anomaly detection models suited to time series data.

Aside from the effectiveness of the methods used, market data are also assumed to carry information that these methods can make use of. It is then worth noting that the Philippine stock market, an emerging market, has been shown to be not weak-form informationally efficient (Aquino, 2002). That is, historical prices carry some information that could be used to detect the above-mentioned market anomalies. It then becomes natural to explore the applicability of machine learning methods for statistical arbitrage detection. While previous studies have shown their effectiveness in detecting anomalies in developed markets (Tallboys et al., 2022), their applicability in emerging markets is yet to be explored. This is where this work comes in: It uses data from the Philippine stock market to determine whether time series anomaly detection approaches are applicable. In doing so, the study could provide investors with tools they can use as they participate in emerging markets.

1.1 Objectives of the Study

Exploring the efficacy of machine learning–based anomaly detection models as statistical arbitrage detection models in the Philippine Stock Exchange (PSE) is the primary motivation for this study. That is, the study aims to determine if anomaly detection can identify statistical arbitrage opportunities. Specifically, the objectives of the study are as follows:

- to detect statistical arbitrage opportunities for selected stock pairs using TadGAN and AER; and
- to compare the performance of the time series anomaly detection approach with benchmark pairs trading strategies (namely distance and cointegration approaches).

1.2 Scope and Limitations

In assessing the applicability of the two aforementioned models, several assumptions for the market are considered. First, this study assumes that fractional shares *can* be traded and that stocks are fully liquid and tradeable daily. Moreover, only the effects of transaction costs are considered. Other costs such as from slippage, interest, and dividends are not considered in the evaluation of the models' performance. Furthermore, only statistical arbitrage is considered, and arbitrage opportunities are detected only for identified pairs of stocks that meet specified conditions. Finally, different market conditions are not noted, as this study focuses on the applicability and efficacy of the models.

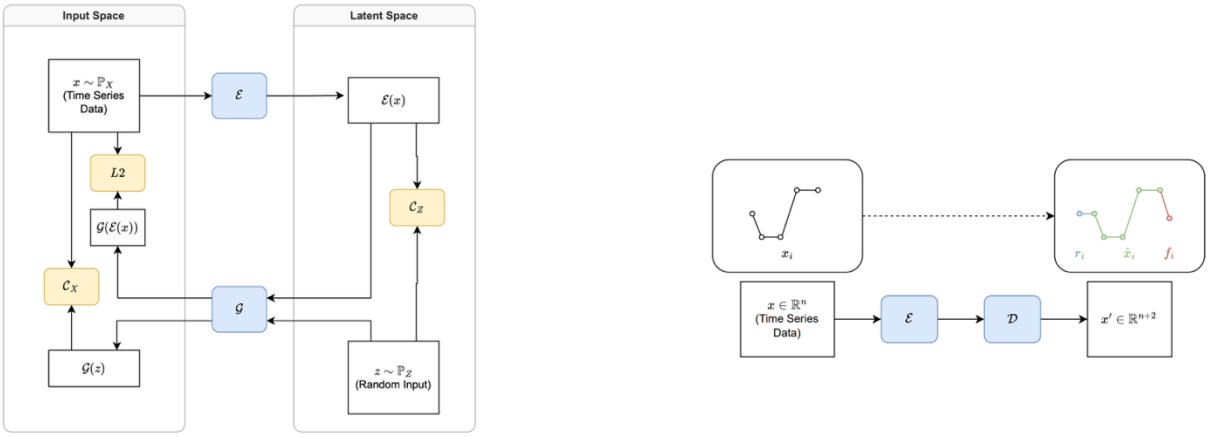
Aside from market assumptions, a few model assumptions are also made. First, the study does not do hyperparameter tuning and optimization. Instead, the models use values based on established practices found in existing related literature. Moreover, in the comparison, the study focusses on simulating and applying the identified pairs trading strategies. Finally, other machine learning models are not considered, as the study is focussed on time series anomaly detection models.

2. Pairs Trading and Anomaly Detection

As mentioned, a commonly used strategy to take advantage of statistical arbitrage involves pairs trading. This strategy relies on monitoring the spread of highly correlated pairs of stocks to generate statistical returns based on a set of systematic rules. In general, pairs trading involves three phases: formation of the pairs, construction of the residual series, and formulating a trading rule. When forming the stock pairs, traders often use statistical properties of stocks as a basis for pairing them. For example, the distance approach utilizes the Euclidean distance of historical prices to identify highly correlated pairs. This distance is monitored, and, when the spread deviates from the historical mean by one standard deviation (Fernández-Pérez et al., 2020), a position is opened. Under the assumption of mean reversion, the spread is expected to return to the historical mean, allowing the position to close at a profit. Although this approach works, using the sum of squared distances is not optimal. Instead, cointegrated pairs of time series are preferred, using a similar method of thresholding to open and close positions. This represents the cointegration approach. Instead of using highly correlated pairs based on their Euclidean distance, the approach uses cointegrated pairs of order 1. These pairs are $I(1)$ processes that have a stationary linear combination, and, by the error correction model, the combination has mean-reverting properties (Rad et al., 2016). However, the problem lies on how long the relationship will continue. Uncertainties on when to close the position may arise when the series does not seem to revert to its historical mean.

Aside from using the spread as a trigger to open positions, positions can also be opened when market anomalies are detected. This is because arbitrage opportunities may be associated with market anomalies. These anomalies are defined as breaks in patterns or trends that are assumed to have been caused by external factors outside of normal market conditions. A common form of anomaly detection is through prediction based on historical values. From this, a confidence interval is found, and when the actual observed value is outside of it, it is labelled an anomaly. An alternative method is through reconstruction. Latent features from the pair's correlation are captured to reconstruct the data instead. When the observed value varies greatly from the reconstructed value, it is labelled an anomaly. New, sophisticated forms of time series reconstruction have been used by incorporating different forms of machine learning methods to improve reconstruction. Latent features from the correlated pairs may be better represented through deep learning models than other models (such as stochastic models). Examples of deep learning models used for anomaly detection are TadGAN (which uses GANs to form time series reconstructions) and AER (which combines anomaly detection from predictive and reconstructive approaches to form a more robust anomaly detection model). These two models are illustrated in Figure 1.

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)



(a) TadGAN (b) AER

Figure 1. Illustration of the architectures of the time series anomaly detection models considered in this paper.

2.1 Time Series Anomaly Detection Using Generative Adversarial Networks (TadGAN)

TadGAN (Geiger et al., 2020) makes use of GANs to reconstruct time series data. It consists of two parts. First is the encoding of the data into a latent space of lower dimension (represented by the mapping function $\mathcal{E}: X \to Z$). This is followed by a generator that decodes information from the latent space reconstructing the time series data (represented by the mapping function $\mathcal{G}: Z \to X$). For each part, a critic is present to act as an adversary to the encoder and generator. In order to avoid the mode collapse problem common to GANs, TadGAN employs Wasserstein loss to train the critics. Furthermore, a cycle consistency loss is used to ensure better reconstruction through the process of encoding and decoding (i.e., $x_i \to \mathcal{E}(x_i) \to \mathcal{G}(\mathcal{E}(x_i))$, where $\mathcal{G}(\mathcal{E}(x_i)) \approx \hat{x}_i$). This gives the following objective function:

$$\min_{\{\mathcal{E},\mathcal{G}\}} \max_{\{\mathcal{C}_x \in \mathbf{C_x}, \mathcal{C}_z \in \mathbf{C_z}\}} V_X(\mathcal{C}_x, \mathcal{G}) + V_Z(\mathcal{C}_z, \mathcal{E}) + V_{L2}(\mathcal{E}, \mathcal{G}). \quad (1)$$

The first two terms are the Wasserstein loss for the critic models given by

$$V_X(\mathcal{C}_x, \mathcal{G}) = \mathbb{E}_{x\sim\mathbb{P}_X}[\mathcal{C}_x(x)] - \mathbb{E}_{z\sim\mathbb{P}_Z}[\mathcal{C}_x(\mathcal{G}(z))], \quad (2)$$
$$V_Z(\mathcal{C}_z, \mathcal{E}) = \mathbb{E}_{z\sim\mathbb{P}_Z}[\mathcal{C}_z(z)] - \mathbb{E}_{x\sim\mathbb{P}_X}[\mathcal{C}_z(\mathcal{E}(x))], \quad (3)$$

where $\mathcal{C}_x(x)$ is the critic score of the original time series x tested against itself, $\mathcal{C}_x(\mathcal{G}(z))$ is the critic score of the series $\mathcal{G}(z)$ generated from random noise tested against the original time series x, $\mathcal{C}_z(z)$ is the critic score of a random latent space sample z tested against itself, and $\mathcal{C}_z(\mathcal{E}(x))$is the critic score of an encoded time series $\mathcal{E}(x)$ tested against a random latent space sample z. Furthermore, note that $\mathbb{P}_X$ and $\mathbb{P}_Z$ are the probability distributions over X and Z, respectively. Thus, $x \sim \mathbb{P}_X$ and $z \sim \mathbb{P}_Z$ indicate that time series x and z are $\mathbb{P}_X$- and $\mathbb{P}_Z$-distributed, respectively. As such, $\mathbb{E}_{x\sim\mathbb{P}_X}[\cdot]$ and $\mathbb{E}_{z\sim\mathbb{P}_Z}[\cdot]$ are the Lebesgue integrals over probability space $\mathbb{P}_X$ and $\mathbb{P}_Z$, respectively. However, these cannot be calculated explicitly and are instead estimated using the arithmetic mean. Finally, the cycle consistency lost based on the L2-norm is given by

$$V_{L2}(\mathcal{E}, \mathcal{G}) = \mathbb{E}_{x\sim\mathbb{P}_X}[x - \mathcal{G}(\mathcal{E}(x))], \quad (4)$$

where x refers to the original time series and $\mathcal{G}(\mathcal{E}(x))$ refers to the reconstructed time series. For the model architecture, the encoder and generator are both built on long short-term memory (LSTM) networks. These models are known for being able to generate outputs based on previously stored information. In contrast, the two critics utilize dense layers with a leaky rectified linear unit (Leaky ReLU) activation function to distinguish between real and generated data.

2.2 Time Series Anomaly Detection Using Auto-Encoder with Regression (AER)

AER, developed by Wong et al. (2022), is another anomaly detection model; it uses both prediction-based and reconstruction-based anomaly detection. These two approaches were considered as they have their respective strengths and weaknesses. It uses an LSTM auto-encoder to reconstruct the time series and an LSTM regressor to predict the subsequent values of the time series. To address the drawbacks of prediction-based anomaly detection, AER's prediction model is bidirectional; it also predicts the previous values of the time series. Consideration of these three parts of the model yields the following objective function:

$$\text{Loss} = \frac{\gamma}{2} V_{\text{pred}}(t_{i-1}, r_{i-1}) + \frac{\gamma}{2} V_{\text{pred}}(t_{i+n}, f_{i+n}) + (1 - \gamma)\, V_{\text{pred}}(t_{i:i+n-1}, y_{i:i+n-1}), \quad (5)$$

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)

where n is the number of observations, γ is the weight of the regression losses on the objective function, $V_{\text{pred}}(a, b)$ is the prediction loss given by the mean squared error between the pairs of values a and b, $V_{\text{rec}}(a, b)$ is the reconstruction error given by the mean squared error between the time series sequences a and b, and t_i, r_i, f_i, and y_i are the actual, reverse predicted, forward predicted, and reconstructed values, respectively. Both the encoder and the decoder utilize bidirectional LSTMs to capture information in both forward and backward directions. The encoder maps the data to a latent space, while the decoder reconstructs the series with two additional data points: one at the beginning for one-step reverse prediction, and one at the end for one-step forward prediction. The rest corresponds to the reconstructed data.

3. Methods and Implementation

To determine the efficacy of TadGAN and AER, post-pandemic daily closing prices from 2021 to 2023 in the Philippine stock market were used. To be included in the dataset, the stock must be among the 100 stocks with the largest market capitalization (valued as of December 2023). They must also be listed in the PSE from 2021 to 2023. In forming the pairs, prices from 2021 to 2022 are taken into account, with the top five pairs being chosen based on the combined criteria of the cointegration and distance methods. The trading period used for simulation entails prices from 2023. Lastly, the pipelines are retrained every 20 trading days. Results from the two anomaly detection models (TadGAN, AER) are compared with benchmark methods (distance, cointegration). A summary of the pairs trading pipeline can be seen in Figure 2. Details of the network architecture for TadGAN and AER used in this paper can be seen in Figure 3.

3.1 Formation of Pairs and Residual Series

The five pairs are chosen using a combination of the distance and cointegration method criteria (Rad et al., 2016). The combined selection criteria have two parts. First, each pair is tested for cointegration of order 1. In this step, the augmented Dickey-Fuller test (Ajewole et al., 2020) is first used to identify which of the stock price time series are $I(1)$ processes. Those that are not $I(1)$ processes are removed from consideration. Afterwards, each pair of $I(1)$ stock prices is tested for cointegration of order 1 using the two-step Engel-Granger method (Moura, 2010). The second part involves computing the sum of squared distances (SSD) for each cointegrated pair found. In this step, the SSD of the stock pair *X* and *Y* is given by $\text{SSD} = \sum_t \left[S_t^{\text{d}}\right]^2$, where

$$S_t^{\text{d}} = \frac{X_t}{X_0} - \frac{Y_t}{Y_0}, \tag{6}$$

and X_t and Y_t are the historical stock prices. The five pairs with the smallest SSD are then chosen. After the pairs are formed, the residual series is constructed. Series S_t^{d} is the residual series used in the distance method; the residual series used in the cointegration method is the stationary linear combination of the cointegrated pair X and Y given by

$$S_t^{\text{c}} = X_t - \gamma Y_t, \tag{7}$$

where γ is the cointegration coefficient obtained from the two-step Engel-Granger method.

3.2 Formation of Trading Rules

The trading rules for the pairs trading strategies considered are as follows. For the distance and cointegration methods, positions are opened when the residual series is far from the "mean" (Fernández-Pérez, 2020). The position is closed upon mean reversion. To be specific, the trading rules for opening a position are the following:

- if the residual series is more than one historical standard deviation *above* the mean, sell X short and buy Y; and
- if the residual series is more than one historical standard deviation *below* the mean, buy X and sell Y short.

While the two methods have similar trading rules, the number of shares involved and the mean used differ. For the distance method, the number of shares involved is one currency unit worth of X and Y, with the mean set to 0. For the cointegration method, however, one share of X and γ shares of Y are used, with the mean set to the historical mean.

For TadGAN and AER, positions are opened when statistical arbitrage opportunities are found. The existence of these opportunities is determined by detecting anomalies in the cointegration residual series. Upon the end of the anomalous sequence, the positions are closed. For these methods, the trading rules are inherited from the cointegration method when the cointegration residuals are used. That is,

- if S_t^{c} is an anomalous point and is *above* the reconstructed series, one share of X will be sold short while γ shares of Y will be bought; and
- if S_t^{c} is an anomalous point and is *below* the reconstructed series, γ shares of Y will be sold short while one share of X will be bought.

3.3 Performance Evaluation

The mean daily per cent profit and loss (%pnl) is used to evaluate the profitability of the aforementioned methods. The formula employed is a generalized version of the one used by Gatev et al. (2006), modified to accommodate the trading rules of the cointegration and machine learning approaches. To be specific, %pnl at time t, denoted by r_t, is

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)

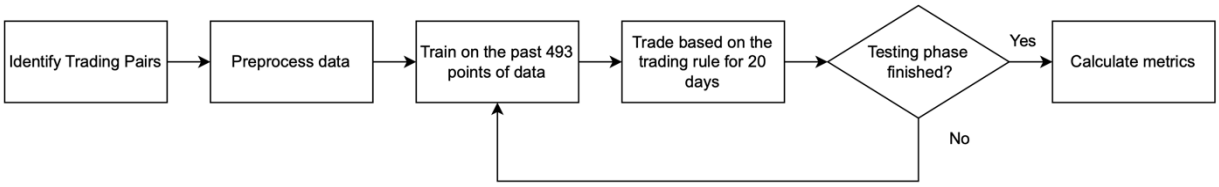


Figure 2. Experiment design flowchart. The pairs trading pipeline includes forming the stock pairs, training the anomaly detection models, and evaluating their profitability.

TadGAN Layer	Output Shape	Activation Function Update	Recurrent
Encoder			
Input	10×1		
LSTM	10×100	Tanh	Sigmoid
Flatten	1000		
Dense	5	Linear	
Critic Z			
Flatten	5		
Dense	100	Linear	
Leaky ReLU	100	Leaky ReLU ($\alpha = 0.2$)	
Dropout	100		
Dense	100	Linear	
Leaky ReLU	100	Leaky ReLU ($\alpha = 0.2$)	
Dropout	100		
Dense	1	Linear	
Generator			
Flatten	5		
Dense	5	Linear	
Reshape	5		
Bidirectional LSTM	5×128	Tanh	Sigmoid
UpSampling1D	10×128		
Bidirectional LSTM	10×128	Tanh	Sigmoid
TimeDistributed Dense	10×1	Linear	
Activation	10×1	Tanh	
Critic X			
Conv1D	8×64		
Leaky ReLU	8×64	Leaky ReLU ($\alpha = 0.2$)	
Dropout	8×64		
Conv1D	6×64		
Leaky ReLU	6×64	Leaky ReLU ($\alpha = 0.2$)	
Dropout	6×64		
Conv1D	4×64		
Leaky ReLU	4×64	Leaky ReLU ($\alpha = 0.2$)	
Dropout	4×64		
Conv1D	2×64		
Leaky ReLU	2×64	Leaky ReLU ($\alpha = 0.2$)	
Dropout	2×64		
Flatten	128		
Dense	1	Linear	

AER Layer	Output Shape	Activation Function Update	Recurrent
Encoder			
Input	10×1		
Bidirectional LSTM	20	Tanh	Sigmoid
Decoder			
Repeat Vector	12×20		
Bidirectional LSTM	12×20	Tanh	Sigmoid
TimeDistributed Dense	12×1		Linear

Figure 3. Breakdown of the layers present in the TadGAN (left) and AER (right) architectures.

$$r_t = \frac{w_{l,t} r_{l,t} - w_{s,t} - r_{s,t}}{C}, \quad (8)$$

where $r_{l,t}$ and $r_{s,t}$ are the percentage increases in the prices of the long and short stocks, respectively. The weights $w_{l,t}$ and $w_{s,t}$ correspond to the values of the respective long and short positions at time t. These weights are indexed by t in order to mimic the evolution of the profit and loss. Finally, C is the value of the long position at the start of the transaction.

To account for transaction costs, these values are adjusted by subtracting a transaction fee as a percentage of the trade. Note that this adjustment is done every time a position is opened or closed. That is, if a transaction occurs at time t, the computed %pnl is adjusted to include a transaction fee. Thus, the adjusted %pnl at time t, denoted by r_t^*, is

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)

$$r_t^* = \begin{cases} r_t - r_{\mathrm{TC}} & \text{if a position is opened/closed at } t \\ r_t & \text{otherwise} \end{cases}, \quad (9)$$

where r_{TC} is the transaction fee expressed as a percentage of the trade. In this work, the transaction fee is assumed to be $r_{\mathrm{TC}} = 0.5\%$.

4. Results and Discussion

4.1 Profit and Loss Analysis

Table 1 shows the unadjusted mean daily %pnl when the following pairs trading strategies are employed for the five stock pairs: TadGAN, AER, the distance approach, and the cointegration approach. With one-sample *t*-test, it was observed that the mean %pnl for the anomaly detection methods is greater than zero for three out of five pairs with at least 90% confidence. Only one pair (PAL–RFM) obtained a positive %pnl with 90% confidence for the cointegration method. Therefore, using TadGAN and AER for pairs trading can yield (more) significant returns than the benchmark methods. This showcases the efficacy of using a time series anomaly detection approach. To be specific, this illustrates the profitability of using machine learning–based anomaly detection models such as TadGAN and AER.

While TadGAN and AER resulted in more profitable trades, some drawbacks worth mentioning were observed. Table 2 shows the number of times a position was opened or closed during the trading period. The anomaly detection methods had more trades than the benchmark methods. This implies that these methods would incur higher transaction costs due to the more frequent buying and selling of stocks. While lower transaction costs for the benchmark methods may appear as an advantage, their infrequent trades (such as the number of trades for the PGOLD–FPH pair using the distance method) might signal another, more pressing drawback of the method, such as the possibility that the trade will never close.

With a hypothetical transaction cost of 0.5% incurred per trade, Table 3 shows that the machine learning methods lost a greater percentage of their profits due to transaction costs. Worth noting is that AER was the most significantly affected method. Despite the loss of profit to transaction costs, AER still outperformed the other three strategies. This can also be observed in Fig. 4 where AER achieves a profit of more than 80% of the total capital employed to all pairs at the beginning of the trading period. However, these results do not account for market liquidity and slippage. Given AER's active trading strategy, trading delays caused by limited trading volume may affect the price trades are executed at, which may reduce the profit generated by the strategy.

4.2 Residual Analysis

Aside from mean daily %pnl and trade counts, residual plots can provide insights into the efficacy of the anomaly detection methods. First, note that for the residual plot of one pair (see Fig. 5a), the position never closed due to the series not crossing 0. That is, mean reversion never happened. This has the potential to incur huge losses due to keeping a short position open. However, looking at the residual plot of another pair (see Fig. 5b), anomaly detection actually improved this by setting a dynamic "normal," as opposed to the static "normal" of the benchmark methods. While this would incur losses, since we expect the movement to be from the residual series rather than the reconstructed series to earn profits, this removes the uncertainty (regarding when to close) that the benchmark methods experience.

While the results presented above indicate the potential of anomaly detection models, it is also important to contextualize them vis-à-vis the nature of emerging markets. For example, emerging markets may harbour characteristics that could affect the profitability of a trading strategy. Compared with developed markets, emerging markets are more volatile. Within emerging markets, it is further observed to be correlated with trading volume (Girard & Biswas, 2007). As such, while increased volatility could provide more pairs trading opportunities, its effects on trading volume may affect its profitability, especially for the anomaly detection strategies.

On the other hand, higher volatility could suggest a more probable occurrence of breakouts. A breakout or breakdown may cause a shift in the trend of a stock's price. This, in turn, may affect the possibility of mean reversion. These phenomena may occur more frequently in emerging markets, making them a concern for the benchmark approaches – these approaches use mean reversion as a condition to close positions. However, the dynamic mean of the anomaly detection approaches addresses this issue. That is, since the anomaly detection approaches do not rely on a static mean, there is no concern that the position will never close.

5. Conclusion

In this paper, we evaluated the efficacy and profitability of using anomaly detection models to identify arbitrage opportunities in an emerging market, namely the Philippine stock market. Different pairs trading strategies – TadGAN, AER, and distance and cointegration approaches – were employed and simulated using historical daily closing stock prices. Results showed that the time series anomaly detection approaches yielded more significant returns as compared to the benchmark approaches. This, however, was at the expense of executing more trades. Nonetheless, the adjusted mean daily %pnl showed that AER still outperformed the other strategies. Furthermore, it should be noted that both TadGAN and AER avoided the problem of positions not being closed due to mean reversion not happening.

It can thus be concluded that TadGAN and AER can be used to detect statistical arbitrage opportunities using the cointegration residuals. That is, anomaly detection models, such as TadGAN and AER, may be applied in pairs trading; of the four strategies tested,

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)

Table 1. Resulting mean daily %pnl for the pairs trading strategies considered. The markers for mean return are interpreted as follows: *** significant at 99% level, ** significant at 95% level, * significant at 90% level.

Pair	TadGAN	AER	Distance	Cointegration
TFHI–STR	0.0062*	0.0090***	0.0033	0.0030
PGOLD–FPH	0.0006	0.0023**	0.0003	0.0001
ALI–MAC	0.0006	0.0016	0.0000	−0.0001
COSCO–FLI	−0.0004	0.0015	−0.0013	−0.0013
PAL–RFM	0.0020**	0.0115***	0.0003	0.0029*

Table 2. Trade counts for each pairs trading strategies considered. It can be seen that strategies employing an anomaly detection approach had more trade counts than the benchmark strategies.

Pair	TadGAN	AER	Distance	Cointegration
TFHI–STR	19	40	7	6
PGOLD–FPH	6	48	3	2
ALI–MAC	11	33	2	2
COSCO–FLI	8	41	1	1
PAL–RFM	8	54	2	8

Table 3. Summary of how aggregate %pnl are affected by transaction costs per trading strategy. A 0.5% fee is subtracted from the %pnl whenever a position is opened/closed.

Strategy	Mean Daily %pnl	Adjusted Mean Daily %pnl	Retained Profit
TadGAN	0.0090	0.0069	76.97%
AER	0.0260	0.0171	66.01%
Distance	0.0026	0.0021	80.19%
Cointegration	0.0046	0.0039	85.14%

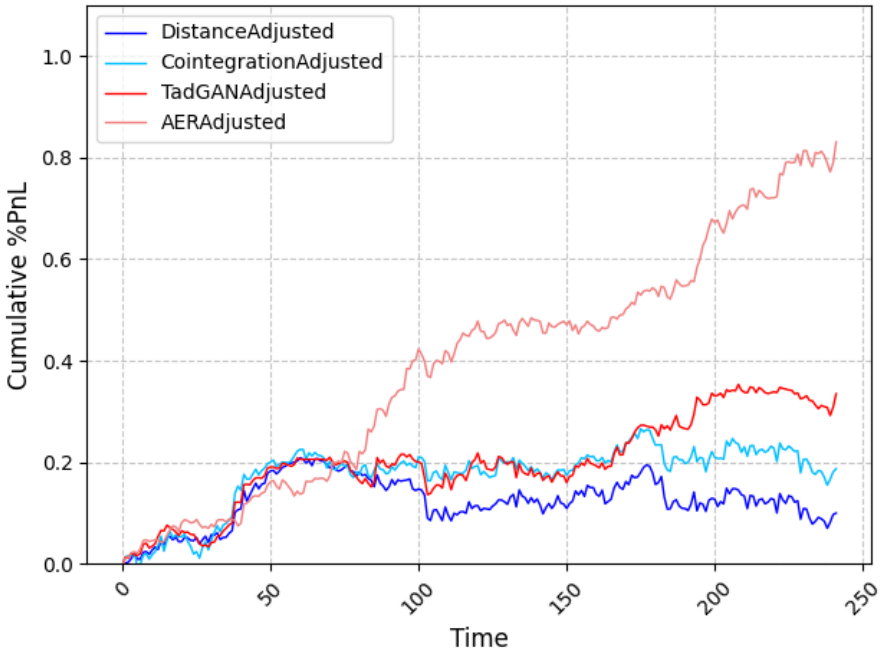


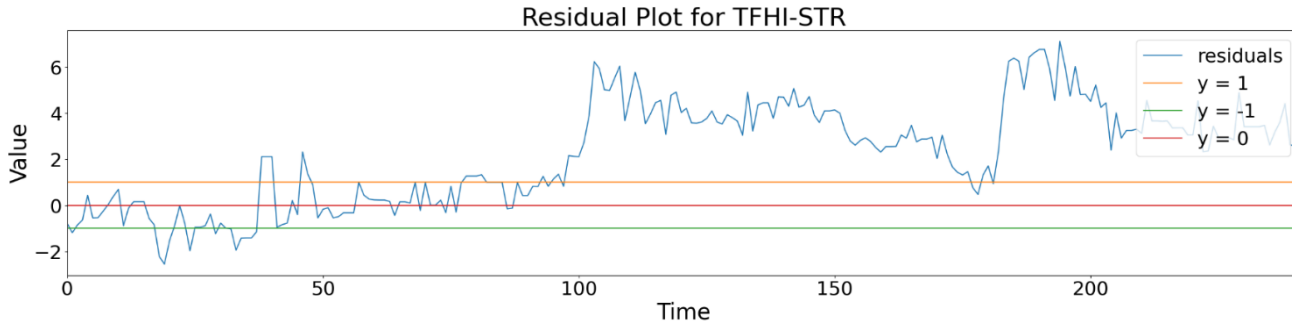
Figure 4. Cumulative %pnl over the trading period for aggregated pairs trading strategies, considering a 0.5% trading cost.

AER had the best performance. The distance and cointegration approaches encountered the problem of positions not closing at mean reversion. This is because at some point, the residual series did not exhibit mean-reverting properties. Contrarily, while the anomaly detection models executed more trades than the benchmark methods (and may close on a loss), overall, they addressed the problem of trades not closing and obtained higher returns than the benchmarks.

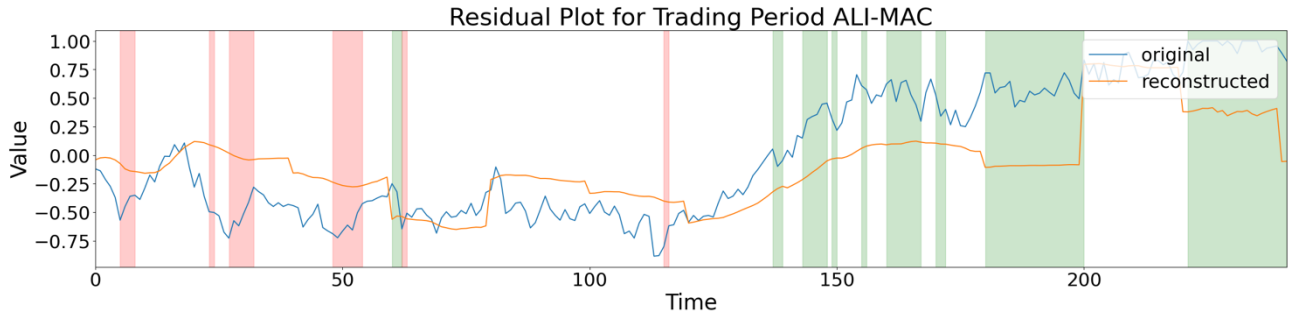
5.1 Future Considerations

While this paper was able to establish the applicability and efficacy of TadGAN and AER in pairs trading, their viability must be further tested. Considering other factors, such as transaction costs, interest expenses, or dividends, and incorporating further data such as bid–ask spreads may provide a more realistic

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)



(a) Cointegration residuals of one stock pair



(b) Original and reconstructed cointegration residuals of another stock pair

Figure 5. Cointegration residuals plots of two stock pairs. The green region indicates when the first trading rule (short the first stock, long the second) is triggered, while the red region is for the second rule (long the first stock, short the second). **(a)** illustrates the basis for how trading is triggered for the benchmark methods. **(b)** illustrates the original and reconstructed residuals, which are used to identify when trading is triggered for the anomaly detection methods.

measure of AER's performance. As the feasibility of the anomaly detection models has been established in this work, their practicability can be considered in future work.

Furthermore, a return calculation tailored to the cointegration approach could be formalized, as this study provided its own performance metric (%pnl). Note that this metric is not exactly similar to returns. Formulating this could more accurately represent the performance of these strategies. Lastly, improvements to the two models can still be explored to enhance their profitability. For example, utilizing both models' residuals could potentially address their limitations and help improve how anomalies are detected. Hyperparameter tuning and optimization could also be considered to optimize the resulting profit and loss metrics.

6. Acknowledgement

The authors would like to thank the Ateneo de Manila University Department of Mathematics and Kristian Mamforte, Christian Paul Chan Shio, and Juan Carlo Mallari for their constructive feedback on this work. A conference presentation of this work was funded by the Ateneo John Gokongwei School of Management Dean's Office.

7. References

Ajewole, K., Adejuwon, S., & Jemilohun, V. (2020). Test for stationarity on inflation rates in Nigeria using augmented Dickey Fuller test and Phillips-Persons test. *IOSR Journal of Mathematics*, *16*(3), 11–14.

Aquino, R. (2002). *Informational efficiency characteristics of the Philippine stock market* (Technical Report). UP College of Business Administration.

Canady, T. S. (2022). *Making sense of market anomalies*.

Carrasco Blázquez, M., De la Orden De la Cruz, C., & Prado Rom´an, C. (2018). Pairs trading techniques: An empirical contrast. *European Research on Management and Business Economics*, *24*(3), 160–167.

Fernández-Pérez, A., de las Nieves López-García, M., & Requena, J. P. R. (2020). On the sensibility of the pairs trading strategy: The case of the FTS Stock Market Index. *Studies of Applied Economics*, *38*(3).

Gatev, E., Goetzmann, W., & Rouwenhorst, K. G. (2006). Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, *19*(3), 797–827.

Geiger, A., Liu, D., Alnegheimish, S., Cuesta-Infante, A., & Veeramachaneni, K. (2020). *TadGAN: Time series anomaly detection using generative adversarial networks*.

Girard, E., & Biswas, R. (2007). Trading volume and market volatility: Developed versus emerging stock markets. *Financial Review*, *42*(3), 429–459.

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)

Krauss, C. (2017). Statistical arbitrage pairs trading strategies: Review and outlook. *Journal of Economic Surveys*, *31*(2), 513–545.

Moura, M. (2010). Testing the Taylor model predictability for exchange rates in Latin America. *Open Economies Review*, *21*, 547–564.

Nakashima, T., & Yairi, T. (2023). Assessing the performance of transformer for time series anomaly detection. In *PHM Society Asia-Pacific Conference* (Vol. 4).

Rad, H., Low, R. K. Y., & Faff, R. (2016). The profitability of pairs trading strategies: Distance, cointegration and copula methods. *Quantitative Finance*, *16*(10), 1541–1558.

Tallboys, J., Zhu, Y., & Rajasegarar, S. (2022). Identification of stock market manipulation with deep learning. In *Advanced Data Mining and Applications* (pp. 408–420). Springer International Publishing.

Ullah, F., Ullah, S., Srivastava, G., & Lin, J. C.-W. (2024). IDS-INT: Intrusion detection system using transformer-based transfer learning for imbalanced network traffic. *Digital Communications and Networks*, *10*(1), 190–204.

Wong, L., Liu, D., Berti-Equille, L., Alnegheimish, S., & Veeramachaneni, K. (2022). *AER: Auto-encoder with regression for time series anomaly detection*.

DOI: https://doi.org/10.22452/mjs.vol45sp2.2
Malaysian Journal of Science 45 (Special Issue 2): 16-24 (November 2026)