

To appear in *Quantitative Finance*, Vol. 00, No. 00, Month 20XX, 1–40

Temporal Representation Learning for Stock Similarities and Its Applications on Investment Management

Yoontae Hwang¹, Stefan Zohren², and Yongjae Lee¹

¹Ulsan National Institute of Science and Technology, Ulsan, Republic of Korea ²University of Oxford, Oxford, UK

(June 01, 2024)

In the era of rapid globalization and digitalization, accurate identification of similar stocks has become increasingly challenging due to the non-stationary nature of financial markets and the ambiguity in conventional regional and sector classifications. To address these challenges, we propose **SimStock**, a novel temporal self-supervised learning framework that combines techniques from self-supervised learning (SSL) and temporal domain generalization to learn robust and informative representations of financial time series data. **SimStock** introduces a dimension corruption method that integrates temporal patterns into the corruption process, enabling the model to learn representations that are resilient to noise and non-stationarity. The primary focus of our study is to understand the similarities between stocks from a broader perspective, considering the complex dynamics of the global financial landscape. We conduct extensive experiments on four real-world benchmarks with thousands of stocks and demonstrate the effectiveness of **SimStock** in finding similar stocks, outperforming existing methods in both same exchange and different exchange scenarios. The practical utility of **SimStock** is showcased through its application to various investment strategies, such as pairs trading, index tracking, and portfolio optimization, where it leads to superior performance compared to conventional methods. Our findings empirically examine the potential of data-driven approach to enhance investment decision-making and risk management practices by leveraging the power of temporal self-supervised learning in the face of the ever-changing global financial landscape.

Keywords: Self-supervised learning, Temporal domain generalization, Financial time series, Stock similarity, Pairs trading, Index tracking, Portfolio optimization

JEL Classification: C45, C58, C67, G11, G17

1. Introduction

The identification of similar stocks is a crucial task in finance, with far-reaching implications for portfolio diversification, risk management, and investment strategy. Traditionally, the similarity between stocks has been determined based on factors such as sector classification, geographical location, and market capitalization. However, in the era of rapid globalization and digitization, these conventional approaches have become increasingly inadequate in capturing the complex dynamics of the global financial landscape. The rapid pace of globalization has led to the emergence of intricate relationships among companies across different sectors and regions. Multinational corporations now operate in multiple countries, blurring the lines between traditional sector classifications.

Corresponding author. Email: stefan.zohren@eng.ox.ac.uk
Corresponding author. Email: yongjaelee@unist.ac.kr

Additionally, the rise of digital technologies has given birth to new industries and business models that transcend geographical boundaries. As a result, the similarities between stocks can no longer be accurately captured by relying solely on regional or sector-based classifications. Moreover, the non-stationary nature of financial markets poses significant challenges to the identification of similar stocks. The statistical properties of financial time series, including returns and correlations, are subject to change over time, a phenomenon known as concept drift or temporal shift (Lu *et al.* 2018, Bai *et al.* 2022). Consequently, similarity measures based on historical data may become unreliable and fail to capture the evolving dynamics of the market.

Recent advances in deep learning-based methods in finance domain have shown promise in predicting the desired parameters from the given data (Zhang *et al.* 2020). However, financial time series data presents unique challenges. The parameters of interest, such as returns, are not only difficult to predict due to the presence of excessive noise, but also the labels that can help the model extract hidden patterns in the data are often not well-defined. Moreover, the complex interactions among various financial assets and the impact of external factors, such as macroeconomic variables and market sentiment, further complicate the estimation process.

One promising solution to tackle these challenges is to leverage representation learning techniques to extract meaningful embeddings from unlabeled financial time series data, taking into account for temporal distribution shifts. By learning robust and informative representations, we can capture the complex relationships between different assets and exploit the inherent structure in the data. Self-supervised learning (SSL) has emerged as a powerful paradigm for learning such representations from unlabeled data, with successful applications in various domains, including computer vision (Chen *et al.* 2020, He *et al.* 2020) and natural language processing (Devlin *et al.* 2018, Yang *et al.* 2019). SSL enables the model to learn meaningful representations by solving pretext tasks that do not require explicit labels, making it particularly suitable for scenarios where labeled data is scarce or expensive to obtain. However, while these methods are suitable for exploiting inductive biases in the data used in any domain, they are not yet suitable for financial time series data.

Also, the application of SSL to financial time series data presents unique challenges. Time series data exhibits distinct characteristics, such as seasonality, trend, and interaction between stocks, which require careful consideration when designing SSL frameworks. Furthermore, the non-stationary nature of financial markets necessitates the development of techniques that can adapt to distribution shifts and generalize well to future time periods. Most existing SSL methods focus on learning invariant representations (Chen *et al.* 2020, He *et al.* 2020), assuming that the data distribution remains stationary. However, this assumption does not hold in the context of financial markets, where the underlying dynamics can change over time.

To address these challenges, we propose a novel methodology called Temporal self-supervised learning, which combines techniques from SSL and the field of temporal domain generalization. Our approach aims to learn general model representations that can adapt to temporal shifts over time, thereby enhancing the accuracy and robustness of financial parameter estimation using learned embeddings. By incorporating temporal domain generalization into the SSL framework, we enable the model to learn representations that are not only informative but also resilient to distribution shifts.

The main contributions of this paper are as follows:

- We propose a novel temporal Self-Supervised Learning framework, **SimStock**, that combines SSL with temporal domain generalization to learn robust and informative representations of financial time series data. **SimStock** leverages the power of SSL to capture the complex relationships between different financial assets while accounting for temporal shifts in the data distribution.
- We introduce a novel corruption method for SSL of stock data, termed dimension corruption, which integrates temporal patterns into the corruption process. By corrupting the input data along different dimensions, such as time, asset, and feature, **SimStock** learns representations that are robust to noise and non-stationarity in the data.

- We conduct extensive experiments on four real-world benchmarks with thousands of stocks to demonstrate the effectiveness of **SimStock** in finding similar stocks. Our results show that **SimStock** achieves state-of-the-art performance, outperforming existing methods in terms of accuracy and robustness. We showcase the practical utility of **SimStock** in simplifying the process of screening potential investment opportunities.
- We demonstrate the practical utility of **SimStock** in various financial applications, including pairs trading, index tracking, and portfolio optimization. We show that using the similar stocks identified by **SimStock** leads to superior performance compared to conventional methods in these applications, highlighting the potential of our approach to enhance investment strategies and risk management.

2. Related work

This section explores relevant research in two key areas: Self-Supervised Learning (SSL) for time-series data and the estimation of financial parameters from historical data. We discuss the challenges and limitations of applying SSL techniques to time series and review various approaches aimed at improving the robustness and accuracy of parameter estimation in finance.

2.1. Self-supervised learning for time-series data

Self-Supervised Learning (SSL) has emerged as a promising paradigm for learning robust and generalizable representations from unlabeled data, with successful applications in both computer vision and natural language processing (NLP) (Chen *et al.* 2020, Jing and Tian 2020, Zhai *et al.* 2019, Lee *et al.* 2019, Qiu *et al.* 2020, Ruder and Plank 2018, Song *et al.* 2020). SSL methods aim to overcome the limitations of traditional supervised learning, which requires large amounts of manually labeled data and can produce models that are brittle and sensitive to small variations in the input.

First and foremost, in the field of NLP, Self-Supervised Learning (SSL) methods, such as BERT (Devlin *et al.* 2018), GPT (Radford *et al.* 2019), and their variants (Mann *et al.* 2020, Touvron *et al.* 2023, Jiang *et al.* 2024), have achieved state-of-the-art results earlier than in other fields by pre-training on large corpora of unlabeled text data to learn to predict masked words or generate realistic text sequences. One reason for the success of SSL in NLP is the abundance of unlabeled text data. Later, SSL techniques were proposed to learn useful visual features by solving pre-tasks on unlabeled images, such as predicting the relative positions of image patches or identifying distorted versions of an image, in order to alleviate the label requirements and enable powerful feature extraction in computer vision as well (Wei *et al.* 2022, Fang *et al.* 2022). This series of processes is not only due to the abundance of unlabeled text data in NLP, but also because language data is discrete and structured (i.e., words), while image data is high-dimensional, continuous, and amorphous, as discussed by (He *et al.* 2020), the success of SSL methodologies in NLP followed by success in computer vision is a natural result.

In the temporal domain, recent research on the SSL method has predominantly concentrated on video understanding (Jenni *et al.* 2020) or action classification (Qian *et al.* 2021). Video-based data is particularly suitable for SSL approaches because of the high correlation between frames, despite its temporal nature. Also, the temporal coherence and continuity present in videos provide a rich source of information that can be exploited to learn meaningful representations without the need for extensive manual labeling. Therefore, methodologies used in computer vision, such as contrastive learning and predictive modeling, can be easily adapted to the video domain. For instance, recent research leverage SSL techniques to learn patch-level (Yun *et al.* 2022, Caron *et al.* 2021) or region-level (Xiao *et al.* 2021) representations from videos. These approaches aim to capture the inherent structure and relationships within video frames, enabling the learning

of meaningful features without relying on explicit labels. By exploiting the temporal coherence and spatial consistency present in videos, these self-supervised methods can discover emergent properties and capture rich semantic information.

However, limited research has been conducted regarding the application of such techniques to financial data reflecting its temporal characteristics. One of the main reasons for this is the challenge of generating different *views* (both positive and negative), which play a key role in self-supervised representation in non-stationary time-series data. In the case of stock data, it is not straightforward to generate meaningful positive and negative views that capture the inherent temporal dependencies and preserve the essential characteristics of the original time series. Unlike images or text, where various augmentation techniques can be applied to create different *views* (e.g., cropping, flipping, masking or word substitution) without significantly altering the core information, financial time series data is highly sensitive to temporal order and contextual relationships. Previous studies have applied various methods to generate views for time series data in order to solve this problem. (Yue *et al.* 2022) utilized timestamp masking and random cropping, while (Choi and Kang 2023) employed scaling and permutation techniques. However, these approaches may not be suitable for time series data that require fine-grained representations.

Next, multivariate time series, which are common in most time series data, present a unique challenge for Self-Supervised Learning (SSL) methods. Although these time series contain multiple dimensions, the relevant and informative patterns are often confined to a limited number of dimensions. Consequently, directly applying SSL techniques that have been successful with other data types to extract meaningful information from time series data becomes difficult, as the sparse nature of useful information in multivariate time series hinders the effectiveness of these methods.

Third, time series data possess distinct characteristics, including seasonality, trend, and frequency domain information (Wen *et al.* 2020, Wu *et al.* 2021a). Considering how SSL works, in sequential data, the hidden patterns of the data are used based on the "current time" that the data is being learned. In other words, this assumes that these characteristics will apply the same in the future. However, the non-stationary nature of time series data poses a significant challenge for SSL methods. The underlying patterns and relationships in the data can change over time, leading to a phenomenon known as concept drift or temporal shift. This means that the representations learned by SSL methods based on historical data may not generalize well to future time periods. To address this issue, techniques from the field of Domain Generalization (DG) can be employed. DG aims to learn general model representations that can adapt to temporal shifts over time. To improve the generalization ability of the model when there is temporal change, methodologies such as DRAIN (Bai *et al.* 2022) have been proposed, but they are only applicable to supervised learning.

2.2. *Parameter estimation in non-stationary financial markets*

In the field of investment management, estimating future returns and risk is a fundamental challenge due to the inherent uncertainty and non-stationary nature of financial markets. Traditionally, researchers and practitioners have relied on historical financial data to estimate various parameters, such as expected returns, volatility, and covariance matrices, which are crucial inputs for portfolio optimization, risk management, and asset pricing models.

One common approach is to use a rolling window of historical returns to estimate the expected returns and covariance matrix of a set of assets. For instance, in the classic mean-variance portfolio optimization (MVO) framework proposed by (Markowitz 1952), the optimal portfolio weights are determined based on the estimated expected returns and covariance matrix using a sample of historical returns. However, an inherent drawback of MVO is the high sensitivity of the optimal portfolio to estimation errors in the input parameters, particularly in the expected returns (Chopra and Ziemba 2013, Chung *et al.* 2022). Small changes in the estimated expected returns can lead to significant shifts in the optimal portfolio weights, resulting in portfolios that may be suboptimal or unstable out-of-sample. Similarly, the Capital Asset Pricing Model (CAPM) (Sharpe 1964, Lintner 1975) and the Fama-French factor models (Fama and French 1993, 2015) rely on historical data to

estimate the beta coefficients and factor premiums.

However, using historical data for parameter estimation has several limitations. First, financial markets are known to exhibit non-stationarity, meaning that the statistical properties of the data, such as the mean and variance, can change over time (Cont 2001, Lo 2017). This implies that the estimates based on historical data may not be representative of the future. Second, the sample size of historical data is often limited, leading to estimation errors and potential over-fitting (Kan and Zhou 2004, 2007).

To mitigate these issues, researchers have proposed various techniques to improve the robustness and accuracy of parameter estimation. One popular approach is the shrinkage method, which combines the sample estimates with a structured estimator to reduce the estimation error (Ledoit and Wolf 2003, 2004b,a). Another technique is the use of robust estimators, such as the minimum covariance determinant (MCD) estimator (Rousseeuw and Driessen 1999) and the minimum volume ellipsoid (MVE) estimator (Van Aelst and Rousseeuw 2009), which are less sensitive to outliers and heavy-tailed distributions. More recently, (Gerber *et al.* 2021) proposed the Gerber statistic for estimating the covariance matrix between assets, a robust co-movement measure that extends Kendall’s Tau by counting the proportion of simultaneous co-movements when their amplitudes exceed data-dependent thresholds, capturing meaningful co-movements while being insensitive to extreme values and noise.

In addition to these techniques, researchers have also explored various approaches to extrapolate past data and model future market dynamics. For example, (Barberis 2000) proposed a Bayesian approach that combines the sample estimates with prior beliefs about the asset returns to improve the out-of-sample performance of portfolio optimization. (Rapach *et al.* 2010) employed combination forecasts, which aggregate individual forecasts based on different predictors, to enhance the accuracy of out-of-sample stock return predictions. Also, (Welch and Goyal 2008) found that the historical average excess return of stocks over bonds, which is often used as an estimate of the equity premium, is sensitive to the choice of the sample period and the assumptions about survivorship bias. Moreover, they showed that this historical average is a poor predictor of future returns, as the equity premium exhibits substantial time-variation and mean-reversion.

Despite these advancements, the fundamental challenge of estimating future returns and risk from historical data remains. The assumption that future behavior will be similar to the past is often violated in practice, as financial markets are subject to regime shifts, structural breaks, and extreme events (Ang and Bekaert 2002, Guidolin and Timmermann 2007).

3. SimStock

We propose `SimStock`, which is graphically illustrated in Figure 1. A key distinguishing characteristic of our framework, setting it apart from prior research, is its utilization of stock augmentation and temporal domain generalization techniques specifically designed to capture the dynamic nature of stock data.

3.1. Preliminary

We consider a self-supervised task where the stock data distribution evolves over time. In the training phase, we are given T observed source domains $D_{1:T} = \{D_1, D_2, \dots, D_T\}$, which are sampled from distributions at T different time points $t_1 \leq t_2 \leq \dots \leq t_T$. Each source domain is denoted as $D_s = \{x_i^s, c_i^s\}_{i=1}^{N_s}$, for $s = 1, 2, \dots, T$, where $x^s \in \mathbb{R}^{d_m}$ represent the d_m -dimensional temporal features, $c^s \in \mathbb{R}^{d_n}$ is d_n -dimensional static metadata, and N_s is the sample size at timestamp t_s . We have omitted the sample index i for simplicity. The model will only be tested on a target domain in the future, i.e., D_{T+1} where $t_{T+1} \geq t_T$.

Our goal is to proactively capture the drift from temporal domains to find stock representations that are robust with respect to temporal distribution shifts. We presume that the representation

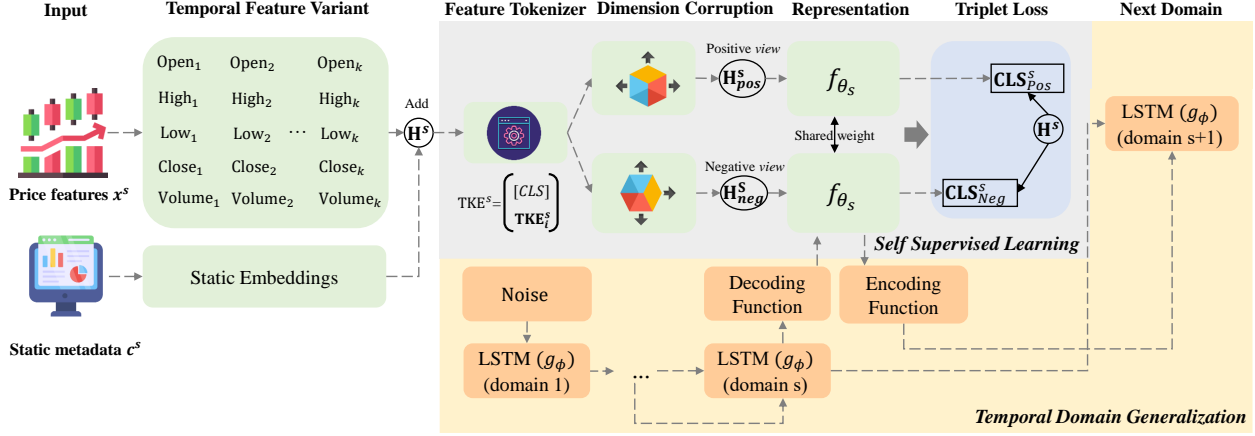


Figure 1. The proposed model (SimStock) combines self-supervised learning framework with temporal domain generalization for stock representations.

model, denoted as f_{θ_s} , is characterized by a deep neural network with function parameters θ_s at timestamp t_s . Consequently, we can get the representation embedding $z_s = f_{\theta_s}(x^s, c^s)$, where z_s represent both temporal and static features of stock data. In the next section, we show that the representation model serves as a mapping function during training, which predicts the dynamics across the parameter $\theta_{1:T} = \{\theta_1, \theta_2, \dots, \theta_T\}$ at each domain D_s .

3.2. Training Dynamics of Temporal Domain Generalization

We are motivated by DRAIN (Bai *et al.* 2022), which first proposed the concept of temporal domain generalization. In each temporal domain D_s , the representation network f_{θ_s} can be trained by maximizing the conditional probability $\mathbb{P}(\theta_s | D_s)$. Here, θ_s signifies the state of the model parameters at timestamp t_s . Given the dynamic nature of D_s , the conditional probability $\mathbb{P}(\theta_s | D_s)$ will also change over time. The objective in the context of temporal domain generalization is to estimate θ_{T+1} utilizing all the training data from $D_{1:T}$. From a probabilistic perspective, we can express this as:

$$\mathbb{P}(\theta_{T+1} | D_{1:T}) = \int_{\Omega} \mathbb{P}(\theta_{T+1} | \theta_{1:T}, D_{1:T}) \cdot \mathbb{P}(\theta_{1:T} | D_{1:T}) d\theta_{1:T}, \quad (1)$$

where Ω denotes the space for model parameters $\theta_{1:T}$. In Eq. 1, the first term inside the integral $\mathbb{P}(\theta_{T+1} | \theta_{1:T}, D_{1:T})$ represents the inference phase, which is the process of predicting the future state of the target representation network (i.e., θ_{T+1}) given all historical states (i.e., $\theta_{1:T}, D_{1:T}$). The second term $\mathbb{P}(\theta_{1:T} | D_{1:T})$ signifies the training phase, which involves leveraging all training data $D_{1:T}$ to ascertain the state of the model on each source domain. More specifically, the training phase can decomposed as follows:

$$\begin{aligned} \mathbb{P}(\theta_{1:T} | D_{1:T}) &= \prod_{s=1}^T \mathbb{P}(\theta_s | \theta_{1:s-1}, D_{1:s}) \\ &= \mathbb{P}(\theta_1 | D_1) \prod_{s=2}^T \mathbb{P}(\theta_s | \theta_{1:s-1}, D_{1:s}). \end{aligned} \quad (2)$$

This formula breaks down the training process into T steps, where the first step corresponds to learning the model parameter on the first domain \mathcal{D}_1 , and each subsequent step corresponds to

learning the model parameter on the new domain \mathcal{D}_s conditional on parameter statuses from the history domains and training data, i.e., $\mathbb{P}(\theta_s \mid \theta_{1:s-1}, \mathcal{D}_{1:s}), \forall s \in \{2, \dots, T\}$.

Suppose that we are at time t_s . In order to effectively address the temporal drift present across the domain, the next parameters θ_{s+1} need to be updated on the current and previous domains $D_{1:s}$. The main problem is how to actually update θ_{s+1} . In this regard, DRAIN introduces a sequential learning process using LSTM (Hochreiter and Schmidhuber 1997) to describe the stochastic process of θ_s . Within the LSTM, each unit g_ϕ defined by its parameters ϕ is used to generate θ_{s+1} while taking into consideration the preceding context $D_{1:s}$ and $\theta_{1:s}$. This process is illustrated with yellow boxes in Figure 1.

3.3. Temporal Representation Learning

Our ultimate goal is to learn a representation model, f_{θ_s} , which captures the stock data distribution that evolves over time. To achieve this, we develop an SSL framework for temporal representation learning of stock data.

Temporal feature variant. Let $x^s \in \mathbb{R}^{d_m}$ be the price feature of a stock, where d_m is the dimension of the feature. The time-varying patterns of stock prices are essential for identifying short- and long-term characteristics of stocks. To learn more rich representations, the price feature x^s is processed by a temporal transformation module μ . Specifically, the price feature x^s is provided with k variations, denoted as:

$$\mu(x^s) = \text{CONCAT}(\mu_1(x^s), \mu_2(x^s), \dots, \mu_k(x^s)) \in \mathbb{R}^{d_{mk}}, \quad (3)$$

where $d_{mk} = d_m \times k$, and each $\mu_i : \mathbb{R}^{d_m} \rightarrow \mathbb{R}^{d_m}$ for $i \in \{1, 2, \dots, k\}$ is a temporal transformation function from the collection U . In other words, $\mu_1, \mu_2, \dots, \mu_k \in U$, where U denotes the collection of temporal transformations. This module is used to create temporal features that incorporate various time intervals. For example, $\mu_1(x^s)$ and $\mu_2(x^s)$ would reflect temporal patterns within a day and a week. Various methods, such as the moving average (Woo *et al.* 2022, Wu *et al.* 2021b), Fourier transform (Zhou *et al.* 2022), and mixtures of experts (Zhou *et al.* 2022), can be utilized to create these temporal features. In this study, we use moving average, which is the most common choice. More specifically, for each temporal transformation μ_i , we calculate the moving average of the price feature x^s over a window of size w_i :

$$\mu_i(x^s) = \frac{1}{w_i} \sum_{j=0}^{w_i-1} x_{t-j}^s, \quad i = 1, 2, \dots, k. \quad (4)$$

where w_i is the window size for the i -th transformation, and x_{t-j}^s denotes the price feature at time $t - j$ in domain s . In our implementation, we use five different window sizes: $w_1 = 5$, $w_2 = 10$, $w_3 = 15$, $w_4 = 20$, and $w_5 = 25$. These window sizes correspond to weekly, bi-weekly, tri-weekly, four-weekly, and monthly moving averages, respectively.

Combined embedding with static metadata. In our framework, static metadata c^s , which can include a variety of data such as firm description, 3-statement financial information, and more, is handled in the static embedding layer. However, for the purpose of this study, we have only included sector information in c^s . As a result, an embedding $\text{Embed}(c^s) \in \mathbb{R}^{d_{mk}}$ is obtained. Next, we create a combined embedding that incorporates both the temporal feature variant $\mu(x^s)$ and the embedded static metadata $\text{Embed}(c^s)$. The resulting combined embedding is denoted as follows:

$$\mathbf{H}^s = \mu(x^s) + \text{Embed}(c^s) \in \mathbb{R}^{d_{mk}}. \quad (5)$$

Feature Tokenizer module. We draw inspiration from the tokenizer approach Gorishniy *et al.*

(2021), which transforms input features into token embeddings to obtain more meaningful representations. This method Gorishniy *et al.* (2021) do not reflect the temporal aspect, but we utilize temporal feature variants in the Feature Tokenizer module to capture the time-varying patterns in the stock price data. The feature-wise token embeddings \mathbf{TKE}_j^s for a given feature index j are computed as follows:

$$\mathbf{TKE}_j^s = b_j^s + \mathbf{H}_j^s W_j^s \quad (6)$$

where $b_j^s \in \mathbb{R}^d$ is the j -th feature bias term and $W_j^s \in \mathbb{R}^d$ is the weight vector for the j -th feature. Consequently, the token embeddings $\mathbf{TKE}^s \in \mathbb{R}^{d_c \times d}$ can be obtained by stacking all of the feature embeddings and adding a special [ST] token, which is known to possess the essence of information after training. This is represented as:

$$\mathbf{TKE}^s = \text{STACK}([\text{ST}], \mathbf{TKE}_1^s, \dots, \mathbf{TKE}_{d_{mk}}^s) \quad (7)$$

where $\mathbb{R}^{d_c \times d} = \mathbb{R}^{(d_{mk}+1) \times d}$ denotes the dimension of the combined token embeddings \mathbf{TKE}^s . Therefore, the feature tokenizer plays a crucial role in learning meaningful temporal representations by transforming combined embedding.

Dimension corruption. In self-supervised learning, the main objective is to learn an embedding space where positive pairs(or views) remain close to each other, while negative pairs(or views) are far apart. When generating views, mixup (Zhang *et al.* 2017) or cutmix (Yun *et al.* 2019) methods are most commonly used. These methods are suitable for invariant augmentation of static data (e.g., images), however, these are not suitable for time-series data (e.g., stocks). We generate views for temporal variants on the same instance, unlike conventional SSL methods that use invariant augmentation by using different instances together. For time-series data, mixing different sequences would ruin the entire temporal structure. Therefore, we propose a dimension corruption method for the augmentation of temporal data.

First, we create positive and negative *views*, \mathbf{H}_{pos}^s and \mathbf{H}_{neg}^s , by randomly shuffling the dimensions within the token embeddings \mathbf{TKE}^s . Here, we define two permutation matrices, \mathbf{P}_{pos}^s and \mathbf{P}_{neg}^s , both of size $d \times d$.¹

$$\begin{aligned} \mathbf{H}_{pos}^s &= \lambda \mathbf{TKE}^s + (1 - \lambda) \mathbf{TKE}^s \mathbf{P}_{pos}^s \\ \mathbf{H}_{neg}^s &= (1 - \lambda) \mathbf{TKE}^s + \lambda \mathbf{TKE}^s \mathbf{P}_{neg}^s \end{aligned} \quad (8)$$

In this case, the formula (7) generates positive and negative views for self-supervised learning. The degree of this perturbation in both views is determined by the mixing parameter λ . With $\lambda > 0.5$, the positive view \mathbf{H}_{pos}^s has minor perturbations, maintaining much of the original token embedding. The negative view \mathbf{H}_{neg}^s is more altered, with greater dimension shuffling, deviating more from the original. More specifically, as shown in Figure 2, the positive view corrupted less variation in temporal features, while the negative view corrupted more variation in temporal features. These views allow the model to consider the chronological order when obtaining embeddings. We set $\lambda = 0.7$ as the default value in this study.

Representation module. The representation module f_{θ_s} aims to characterize the shift between different domains by refining the parameters θ_s through the process described in Section 3.2. In order to effectively reflect temporal patterns of corrupted token embeddings (\mathbf{H}_{pos}^s and \mathbf{H}_{neg}^s), we use the self-attention mechanism (Vaswani *et al.* 2017). The self-attention mechanism aggregates

¹A permutation matrix is a square 0-1 matrix that has exactly one entry of 1 in each row and each column and 0s elsewhere.

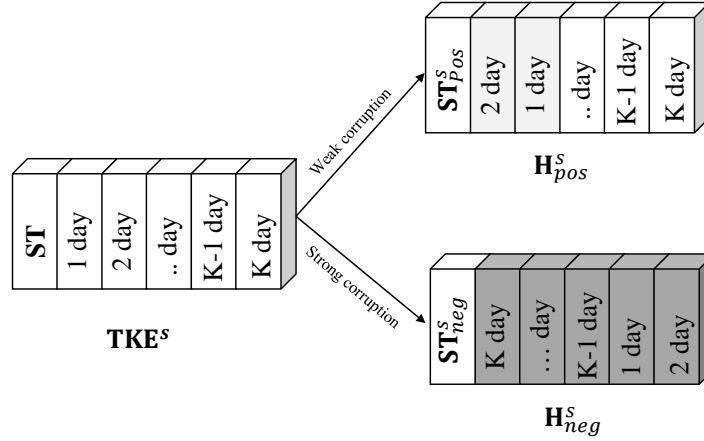


Figure 2. Dimension corruption method for generating positive and negative views from token embeddings (\mathbf{TKE}^s). The positive view \mathbf{H}_{pos}^s is created by applying a small perturbation in dimension order to the original TKE, while the negative view \mathbf{H}_{neg}^s is generated with a larger perturbation in dimension order, preserving more of the original temporal structure in the positive view compared to the negative view.

corrupted token embeddings with normalized importance as follows:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V. \tag{9}$$

Here, $Q = \mathbf{H}_*^s W_Q \in \mathbb{R}^{d \times d_k}$, $K = \mathbf{H}_*^s W_K \in \mathbb{R}^{d \times d_k}$ and $V = \mathbf{H}_*^s W_V \in \mathbb{R}^{d \times d_v}$ represent queries, keys, and values, respectively. Note that \mathbf{H}_*^s represents token embeddings (either positive or negative), while W_Q , W_K , and W_V are learnable matrices that share weights between positive and negative token embeddings. The output, which has a dimension of d_v , is then transformed back into an embedding of dimension d through a fully connected layer. Finally, the outputs \mathbf{ST}_{neg}^s and \mathbf{ST}_{pos}^s are obtained.

Triplet loss. For SimStock, we train it to minimize a triplet loss (Balntas *et al.* 2016), which is a popular choice in SSL. The key idea behind triplet loss is the use of triplets, each of which consists of an anchor, and positive and negative views. Here, the anchor is the embeddings for the unperturbed combined embedding.

For the triplet $(\mathbf{ST}_{pos}^s, \mathbf{ST}_{neg}^s, \mathbf{H}^s)$, where \mathbf{ST}_{pos}^s is the positive view, \mathbf{ST}_{neg}^s is the negative view, and \mathbf{H}^s is the combined embedding (anchor), the triplet loss is defined as follows:

$$L_{\text{triplet}} = \max(0, \text{sim}(\mathbf{H}^s, \mathbf{ST}_{pos}^s) - \text{sim}(\mathbf{H}^s, \mathbf{ST}_{neg}^s) + \alpha) \tag{10}$$

In the above equation, $\text{sim}(\cdot, \cdot)$ denotes a similarity measure (e.g., cosine similarity or Euclidean distance), and $\alpha > 0$ is a margin that is introduced to separate positive pairs from negative pairs. The intuition behind this loss function is that we want to ensure that the anchor point gets closer to the positive sample than to the negative sample by at least the margin α .

Inference phase. In our framework, the inference phase is particularly important. Unlike most existing contrastive representation learning studies (Chen *et al.* 2020, Grill *et al.* 2020), our model is specifically designed to be robust with respect to temporal distribution shifts. The inference phase consists of passing the target domain D_{s+1} through the embedding module to obtain the combined embedding \mathbf{H}^{s+1} and further processed by the feature tokenizer module to obtain the token embeddings \mathbf{TKE}^{s+1} . The stock representation is then obtained by feeding \mathbf{TKE}^{s+1} into the representation model $f_{\theta_{s+1}}$, which is updated with the optimal parameters θ_{s+1} generated by the TDG method described in Section 3.2

4. Experiment

Now we present experiment results to thoroughly demonstrate the performance of `SimStock` on real-world benchmark datasets. The source code is available at <https://github.com/Yoontae6719/SimStock-Representation-Model-for-Stock-Similarities>

4.1. Implementation details

We present the details of tasks, datasets, baseline models, hyperparameter selection, evaluation metrics, and the experiment setting.

4.1.1. Tasks. We conduct four main experiments to assess the effectiveness of `SimStock` in various financial applications:

- **Finding similar stocks:** We evaluate how well `SimStock` can identify stocks that exhibit similar price movements, compared to baseline methods. This is tested in both same exchange (finding similar stocks within the same exchange) and different exchanges (finding similar stocks in a different exchange from the query stock) scenarios.
- **Pairs trading:** We evaluate the effectiveness of a pairs trading strategy based on the similar stocks identified by `SimStock` and baseline models. We form pairs of stocks that are found to be similar, and then execute a pairs trading strategy.
- **Index tracking:** Given an ETF as the query, we find a basket of similar stocks using `SimStock` and assess the performance of using these stocks for index tracking. The goal is to see if the selected stocks can closely mimic the returns of the target ETF.
- **Portfolio optimization:** We investigate whether `SimStock` embeddings can enhance portfolio optimization. Specifically, we replace the correlation matrix of stock returns using the `SimStock` embedding as a similarity measure for mean-variance portfolio optimization. We compare the performance of the resulting portfolios with those obtained using conventional methods for estimating.

4.1.2. Datasets. We collected the daily stock price features (Open, High, Low, Close, and Volume) and sector information for stocks listed on the **NYSE** (New York Stock Exchange), **NASDAQ** (National Association of Securities Dealers Automated Quotations), **SSE** (Shanghai Stock Exchange), **SZSE** (Shenzhen Stock Exchange), and **TSE** (Tokyo Stock Exchange) from Yahoo Finance.

We generate normalized input features describing the trend of a stock on day t . z_{Open} , z_{High} and z_{Low} represent the comparison values of the opening, highest, and lowest prices, respectively, relative to the closing price of the same day. Also, z_{Close} and z_{Volume} represent the comparative values of the closing prices and the volume values compared with day $t-1$, respectively. Refer to Table 1 for the formulas used to calculate each feature. In addition, we calculated stock price features for 5, 10, 15, 20, 25, and 30-day intervals for the temporal feature variant as discussed in 3.3.

Price features	Description
z_{Open}	$\text{Open}_t / \text{Close}_t - 1$
z_{High}	$\text{High}_t / \text{Close}_t - 1$
z_{Low}	$\text{Low}_t / \text{Close}_t - 1$
z_{Close}	$\text{Close}_t / \text{Close}_{t-1} - 1$
z_{Volume}	$\text{Volume}_t / \text{Volume}_{t-1} - 1$

Table 1. Normalized temporal price features.

We use all stocks listed on the **NYSE** and **NASDAQ** (4,231 stocks) and refer to them as the **US** exchange in our study. For the **SSE** (1,407 stocks), **SZSE** (1,696 stocks), and **TSE** (3,882

stocks), we treat each exchange separately. The training period for all tasks is from January 1, 2018, to December 31, 2021, while the test period spans from January 1, 2023, to December 31, 2023. The reference period for comparison to the test period is January 1, 2022 through December 31, 2022. In other words, we verify whether the stocks identified as similar in the reference set maintain their similarity in the subsequent one-year test period.

4.1.3. Baseline models. For the task of finding similar stocks, index tracking and pairs trading task, we compare `SimStock` with the following baselines:

- **Corr1:** Calculates the correlation of stock returns using the past one-year returns (i.e., from January 1, 2022 to December 31, 2022).
- **Corr2:** Calculates the correlation of stock returns using returns from the beginning of the test period (i.e., from January 1, 2018 to December 31, 2022).
- **Peer:** Uses the list of similar stocks provided by Financial Modeling Prep.¹
- **TS2VEC** (Yue *et al.* 2022): A state-of-the-art method based on self-supervised learning for finding similar time-series data.

For the portfolio optimization task, we compare the performance of portfolios constructed using the `SimStock` similarity matrix with those constructed using the following conventional correlation matrices:

- **Historical covariance matrix (HC)** (Jobson and Korkie 1980): The sample correlation matrix of the most recent past stock returns.
- **Shrinkage method (SM)** (Ledoit and Wolf 2004a): A shrinkage estimator of the covariance matrix proposed by (Ledoit and Wolf 2004a).
- **Gerber statistic (GS)** (Gerber *et al.* 2021): A robust correlation measure that counts the proportion of simultaneous co-movements between assets when their amplitudes exceed data-dependent thresholds.

4.2. Can `SimStock` find similar stocks?

In this section, we explore two distinct scenarios for identifying similar stocks using our proposed model: the *same exchange scenario* and the *different exchanges scenario*. The same exchange scenario focuses on finding similar stocks within the same exchange given a query stock. This approach allows for the identification of stocks with comparable characteristics and behaviors within a specific market. On the other hand, the different exchanges scenario involves finding similar stocks within another exchange given a query stock. This scenario leverages the concept of transfer learning, where the trained weights of a model from one exchange are applied to stock data from a different exchange. Transfer learning is a machine learning technique that leverages the knowledge gained from a model trained on a source task to improve the performance of a model on a related target task (Pan and Yang 2009). In our study, we apply transfer learning by utilizing a model trained on one stock exchange (source domain) to identify similar stocks in another exchange (target domain). For example, models trained on the US exchange can be used to find similar stocks in the SSE, SZSE, or TSE exchanges. Notice that the query is not restricted to individual stocks. It can be either sector indices or ETFs.

4.2.1. Evaluation metrics. To evaluate the performance of models in finding similar stocks compared to other models, we employ two widely used metrics: Correlation and Dynamic Time Warping (DTW). DTW is a measure for measuring the similarity between two temporal sequences,

¹<https://site.financialmodelingprep.com/>

allowing for non-linear alignments and capturing the similarities in the overall patterns of stock price movements (Berndt and Clifford 1994). Consider two time-series sequences $X = \{x_1, \dots, x_m\}$ and $Y = \{y_1, \dots, y_n\}$. The DTW between X and Y is defined as:

$$\text{DTW}(X, Y) = \sqrt{\sum_{(i,j) \in \pi} \|x_i - y_j\|^2}. \quad (11)$$

Here, an alignment path π of length K is a sequence of K index pairs $(i, j)_K$, where $\max(m, n) \leq K \leq m + n - 1$. Also, $\|\cdot\|$ is the Euclidean distance. DTW uses global path constraints while comparing two time-series sequences X and Y . That is, the pairs i and j are constrained so that $|i - j| \leq r$, where r is a predefined radius, in the case of the Sakoe–Chiba band.

To assess the effectiveness of our approach in capturing similarities and predicting future stock behaviors, we evaluate the model’s performance using an out-of-sample period. The dataset used for inference is the year immediately preceding the out-of-sample period. By finding similar stocks in the previous year and comparing their performance to the target stock in the out-of-sample period, we measure the model’s performance using Correlation and Dynamic Time Warping (DTW) measures. We consider the top 9, 7, 5, 3, and 1 similar stocks (namely, TOP@9, TOP@7, TOP@5, TOP@3, and TOP@1) for evaluation.

4.2.2. Quantitative evaluation: same exchange scenario. The diagonal plots in Figure 3 and Figure 4 illustrate the performance (Correlation and DTW) of different models in same exchange scenario. The diagonal plots show that our proposed **SimStock** consistently outperforms the baselines (Corr1, Corr2, Peer, and TS2VEC) in finding similar stocks within the same exchange, across all four stock markets (US, SSE, SZSE, and TSE). For example, in the US to US scenario, *SimStock* achieves a correlation of around 0.8 for the top similar stock (TOP@1), substantially higher than the 0.6 correlation of the best performing TS2VEC.

Similarly, for the SSE to SSE and TSE to TSE scenarios, **SimStock** consistently shows superior performance compared to the baselines, with correlations close to 0.8 for TOP@1, but **SimStock** and TS2VEC gave equivalent results in SZSE to SZSE. Note that the performances of all the baseline models were not much different. It is interesting that the peer stocks picked by investment platforms (Peer) were not quite close to the query stocks in terms of DTW. The DTW results also indicate that **SimStock** identifies more similar stocks, as evidenced by the lower DTW distances compared to the baselines. These findings demonstrate that our **SimStock** effectively captures the underlying patterns and dynamics of stock price movements, enabling it to find highly similar stocks within the same exchange.

4.2.3. Quantitative evaluation: different exchanges scenario. The off-diagonal plots in Figure 3 and Figure 4 represents the outcomes of identifying for similar stocks in exchanges different from the exchange of the query stock. Note that Peer is not available for this scenario, because most trading platforms do not provide information on similar stocks in other exchanges. The off-diagonal plots demonstrate that **SimStock** outperforms the baselines in all different exchanges scenarios, indicating its effectiveness in leveraging temporal domain generalization to identify similar stocks across different markets. For instance, in the US to SSE, US to SZSE and US to TSE scenarios, **SimStock** achieves correlations around 0.3 for TOP@1, surpassing the performance of Corr1, Corr2, and TS2VEC. Similarly, for the SSE to US, SSE to SZSE and SSE to TSE scenarios, **SimStock** exhibits higher correlations compared to the baselines, showcasing its ability to capture similarities between stocks in these closely related exchanges.

However, it is important to note that the overall performance of all models in the different exchanges scenarios is lower compared to the same exchange scenario. This can be attributed to the inherent differences in market dynamics, regulations, and economic factors across different ex-

changes. Despite these challenges, **SimStock** consistently outperforms the baselines in all cases, highlighting its robustness and adaptability in finding similar stocks across diverse market conditions.

The DTW results in the off-diagonal plots further support the superiority of **SimStock** in identifying similar stocks across exchanges. The DTW approach is a more suitable alternative to the correlation metric when considering different timestamps. **SimStock** consistently achieves lower DTW distances compared to the baselines, indicating its ability to capture the overall patterns and similarities in stock price movements, even when applied to different markets. These findings underline the potential of **SimStock** in enabling investors and analysts to identify similar stocks across different exchanges.

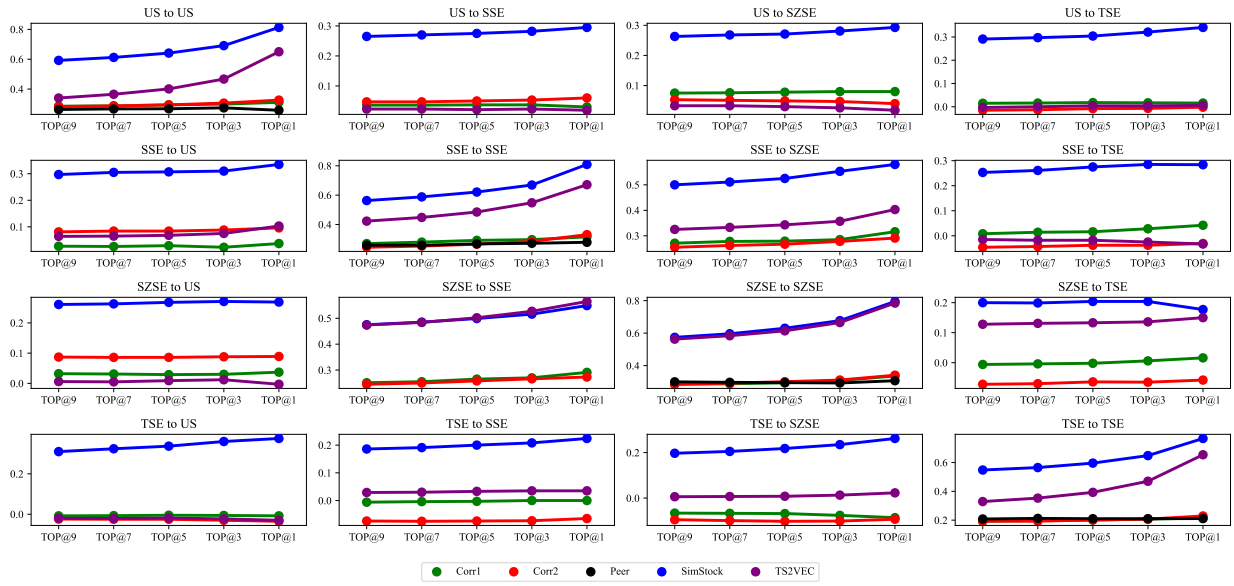


Figure 3. Performance of models in same exchange (diagonal) and different exchanges (off-diagonal) scenarios for finding similar stocks. The performance is evaluated using TOP@ k Correlation metrics, where $k = 9, 7, 5, 3,$ and 1 . Each data point represents the average correlation between the target stock and the top k similar stocks identified by the respective model.

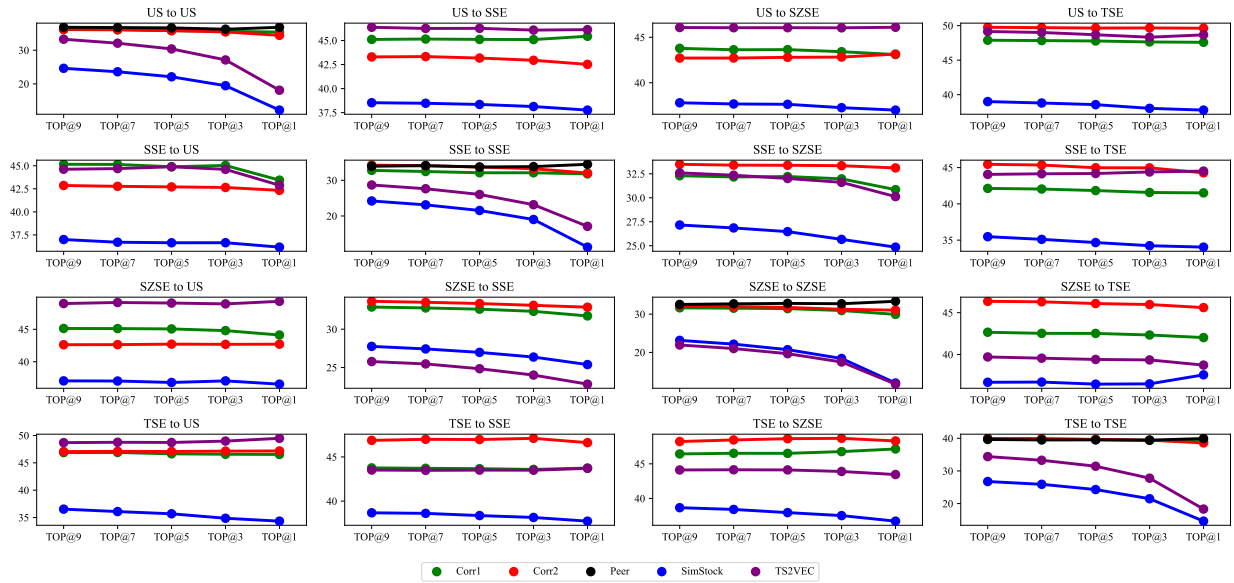


Figure 4. Performance of models in same exchange (diagonal) and different exchanges (off-diagonal) scenarios for finding similar stocks. The performance is evaluated using TOP@ k DTW metrics, where $k = 9, 7, 5, 3,$ and 1 . Each data point represents the average correlation between the target stock and the top k similar stocks identified by the respective model.

4.2.4. Qualitative evaluation. The similarity between stocks should not be measured only on time-series distances. Similar stocks may be in different industries, but they should have similar fundamental information. Unfortunately, however, it is almost impossible to measure such similarity in a quantitative way. Therefore, we provide qualitative analysis of the results.

To assess the effectiveness of SimStock in identifying fundamentally similar stocks, we conduct a nuanced analysis by examining the actual stocks recommended by each method for a diverse set of query stocks from various sectors (Table 2). We compare the top@3 similar stocks identified by SimStock with those of the baseline methods, including TS2VEC, Corr1, Corr2, and Peer, to evaluate whether SimStock captures relevant business characteristics, financial attributes, and industry-specific similarities better than the baselines.

When we begin by examining the technology sector, SimStock identifies highly relevant similar stocks for query stocks such as Apple (AAPL), Microsoft (MSFT), and Meta (META). For AAPL, SimStock's top@ k similar stocks include Microsoft (MSFT), Tyler Technologies (TYL), and Intuit (INTU), all well-established companies in the software and technology domain. In contrast, the baseline methods identify a mix of technology and non-technology companies, such as Amazon (AMZN) for TS2VEC and Corr2, and Thermo Fisher Scientific (TMO) for Corr1. Similarly, for MSFT, SimStock identifies Cadence Design Systems (CDNS), Manhattan Associates (MANH), and Tyler Technologies (TYL) as similar stocks, while the baselines include manufacturing companies like Danaher (DHR) for Corr1.

Next, we analyze the healthcare sector, where SimStock identifies relevant similar stocks for query stocks such as Pfizer (PFE) and Centene (CNC). For PFE, SimStock's top similar stocks include BioNTech (BNTX), Moderna (MRNA), and Johnson & Johnson (JNJ), all major players in the pharmaceutical and biotechnology industries. The baseline methods, however, identify less relevant stocks, such as NantKwest (NKNG) for TS2VEC and Metrocity Bankshares (MCBS) for Corr1. For CNC, SimStock identifies other health insurance and financial services companies like Bank of Montreal (BMO) and Morgan Stanley (MS), while the baselines include less related stocks, such as Humana (HUM) for TS2VEC and Stericycle (SRCL) for Corr2.

In the energy sector, SimStock identifies industry-specific similar stocks for the query stock Exxon-Mobil (XOM), such as Marathon Oil (MRO), Cenovus Energy (CVE), and Hess Corporation (HES). The baseline methods also identify some relevant energy stocks, such as MRO for TS2VEC and Murphy Oil Corporation (MUR) for Corr1 and Marathon Petroleum (MPC) for Corr2. That is, all methods successfully identified stocks that are similar to XOM.

Finally, we examine the financial sector, where SimStock identifies relevant similar stocks for query stocks such as Wells Fargo (WFC) and Visa (V). For WFC, SimStock's top similar stocks include Bank of America (BAC), Fifth Third Bancorp (FITB), and F.N.B. Corporation (FNB), all well-known banking institutions. Similar, the baseline methods also identify relevant banking stocks, such as BAC for Corr2, JPMorgan Chase (JPM) for TS2VEC and Berkshire Hills Bancorp (BHLH) for Corr1. For V, SimStock identifies Mastercard (MA), Stifel Financial (SF), and Intercontinental Hotels Group (IHG) as similar stocks, while the baselines include less related stocks, such as Planet Fitness (PLNT) for Corr2.

4.3. Application to Pairs trading

In this section, we investigate the practical application of the similar stocks identified by SimStock in the context of pairs trading. Pairs trading is a market-neutral trading strategy that exploits the relative mispricing between two highly correlated securities (Gatev *et al.* 2006). The strategy involves simultaneously buying the relatively underpriced security and selling the relatively overpriced security, with the expectation that the prices will converge to their long-term equilibrium. We assess the profitability of pairs trading using the top@ k similar stocks identified by SimStock and compare its performance with benchmark models, namely TS2VEC, Corr1, and Corr2. We added Cointegration method (Coint)(Nelson and Plosser 1982, Enders 2004) to the benchmark model, which selected 3 stocks from the US exchange that exhibited the strongest cointegration

Query Stock	Method	TOP@3 similar stocks			Query Stock	Method	TOP@3 similar stocks		
		First	Second	Third			First	Second	Third
AAPL	SimStock	MSFT	TYL	INTU	PFE	SimStock	BNTX	MRNA	JNJ
	TS2VEC	AMZN	WTM	AMD		TS2VEC	NKNG	DJCO	ESGR
	Corr1	TMO	SNPS	CNDS		Corr1	ICL	MCBS	PTSI
	Corr2	GLOB	AMZN	TYL		Corr2	RMR	NVS	BUD
	Peer	MSFT	NVDA	ASML		Peer	LLY	ABBV	NVO
CMG	SimStock	AMZN	MANH	MSFT	AMZN	SimStock	CMG	INTU	MANH
	TS2VEC	NVR	USLM	NEU		TS2VEC	AAPL	F	PLPC
	Corr1	DHR	DSGX	PCTY		Corr1	ACMR	ADBE	FIVN
	Corr2	PAYC	PCTY	MANH		Corr2	TEAM	LYV	GENE
	Peer	HLT	RACE	AZO		Peer	TSLA	BKNG	SBUX
MSFT	SimStock	CDNS	MANH	TYL	BA	SimStock	IVZ	SPR	UAA
	TS2VEC	GOOG	GOOGL	MA		TS2VEC	LPL	NNI	FCX
	Corr1	DHR	TMO	FAST		Corr1	NCLH	SPR	SOHO
	Corr2	GOOGL	GOOGL	DAVA		Corr2	RVSB	ENVA	STT
	Peer	AAPL	NVDA	ASML		Peer	NOC	CNI	WM
WFC	SimStock	BAC	FITB	FNB	META	SimStock	SPOT	PYPL	FORM
	TS2VEC	JPM	C	MA		TS2VEC	UHAL	MAR	MSFT
	Corr1	BHLH	WNEB	RVSB		Corr1	CVNA	GREE	INDP
	Corr2	BAC	WBS	CFG		Corr2	SKYW	JAGX	CTHR
	Peer	CHTR	NTES	ATVI		Peer	MCD	LOW	TM
V	SimStock	MA	SF	IHG	MA	SimStock	V	BKNG	IHG
	TS2VEC	MA	MSFT	KO		TS2VEC	V	KO	NUE
	Corr1	MA	TDY	ROP		Corr1	V	TDY	FICO
	Corr2	MA	PLNT	RELX		Corr2	V	GES	RTO
	Peer	MA	ADBE	CSCO		Peer	JPM	BAC	V
XOM	SimStock	MRO	CVB	HES	CVS	SimStock	CNC	BMO	MS
	TS2VEC	MRO	CVX	NUE		TS2VEC	HUM	VNR	BKNG
	Corr1	MUR	MRO	EOG		Corr1	CCB	CHRD	RJF
	Corr2	MPC	HES	ERF		Corr2	SRCL	MLM	NMFC
	Peer	CVX	SHEL	TTE		Peer	ANTM	MDT	GSK

Table 2. Top@3 similar stocks identified by SimStock and baseline methods (TS2VEC, Corr1, Corr2, and Peer) for a diverse set of query stocks from the technology, healthcare, energy, and financial sectors.

relationships, as determined by their p -values in descending order.

4.3.1. Pairs trading procedure. We employ the price ratio approach for pairs trading, which involves constructing a spread between two securities by computing the ratio of their prices. The spread is then normalized by its historical mean and standard deviation to obtain the Z-score, which measures the deviation of the current spread from its historical average in terms of standard deviations. The Z-score is calculated as follows:

$$PR_t = P_{q,t}/P_{s,t}$$

$$Z_t = (\mu_1^{PR_t} - \mu_2^{PR_t})/\sigma_2^{PR_t} \quad (12)$$

where $P_{q,t}$ and $P_{s,t}$ are the prices of the query stock and similar stock in the pair at time t , and μ_1 , μ_2 , and σ_2 are the rolling mean and standard deviation of the price ratio up to time t , calculated using lookback periods of L_1 and L_2 trading days, respectively.

Trading signals are generated based on the Z-score, with the opening of a long-short position when the Z-score exceeds a pre-specified threshold of ± 1.25 SD, indicating a significant deviation from the long-term equilibrium. Specifically, when $Z_t > 1.25$ SD, we sell the overpriced security and buy the underpriced security. Conversely, when $Z_t < -1.25$ SD, we buy the overpriced security and sell the underpriced security. The position is closed when the Z-score reverts back within the range $[-0.5$ SD, 0.5 SD], suggesting price convergence. To limit potential losses, a stop-loss (Hwang *et al.* 2023) mechanism is implemented, whereby positions are closed if the portfolio value drops below a predetermined threshold L_{stop} of 500 USD.

To determine optimal values for the hyperparameters L_1 and L_2 , we perform a grid search over the training period from 2022-01-01 to 2022-12-31. Different L_1 and L_2 values are determined for each query stock paired with each of its top 3 similar stocks identified by SimStock and the benchmark models (TS2VEC, Corr1, Corr2), resulting in 3 pairs per query stock per model. Pairs trading is applied independently to each pair, with overall performance for each query stock and model calculated as the average wealth across the 3 pairs, starting with an initial trading capital of 10,000 USD.

4.3.2. Evaluation of pairs trading performance. Table 3 presents the average terminal wealth and maximum drawdown (MDD) achieved by applying pairs trading to the top@3 similar stocks identified by SimStock and the benchmark models (TS2VEC, Corr1, Corr2, Coint) for each query stock. The results of pair trading between query stock and TOP@1 stock can be found in Appendix B.

SimStock achieves the highest average terminal wealth for 6 out of the 12 query stocks (AAP, CMG, MSFT, WFC, BA and CVS). Coint performs best for 2 query stocks (XOM and PFE), while Corr1, Corr2 and TS2VEC each perform best for 1 query stock (MA, AMZN and V respectively). For the remaining query stocks, SimStock ranks second in terms of wealth for 4 cases (V, XOM, META and MA).

Regarding maximum drawdown, SimStock exhibits the lowest MDD for 7 out of the 12 query stocks (AAPL, CMG, MSFT, WFC, PFE, META and MA), indicating better risk management. TS2VEC, Corr1, Corr2 and Coint achieve the lowest MDD for 1 query stock each (AMZN, V, CVS and XOM respectively). For the remaining 5 query stocks (V, XOM, AMZN, BA and CVS), SimStock has the second-lowest MDD. The ability to generate higher returns while maintaining lower drawdowns suggests the robustness of the pairs identified by SimStock.

It is important to note that TS2VEC, Corr2, and Coint failed to generate buy/sell signals for certain stocks. Specifically, TS2VEC and Coint were unable to generate signals for AAPL and WFC, while Corr2 failed to do so for WFC. SimStock, on the other hand, generated a buy/sell signal for one out of the three stocks, specifically META.

The results demonstrate the superior performance of SimStock in identifying profitable pairs for trading compared to the benchmark models. SimStock's ability to achieve the highest average terminal wealth for the majority of query stocks while maintaining the lowest maximum drawdown for most cases highlights its effectiveness in capturing meaningful similarities between stocks that translate into successful pairs trading strategies.

Query Stock	Wealth				
	SimStock	TS2VEC	Corr1	Corr2	Coint
AAPL	961.04 ±474.43	NaN**	234.69 ±1165.48	916.07 ±338.95	NaN**
CMG	546.95 ±724.24	282.38 ±714.16	-1070.09 ±722.33	-1098.98 ±893.35	-857.35 ±2967.72
MSFT	754.12 ±69.73	498.11 ±785.76	474.85 ±1651.85	-306.61 ±2114.7	257.29 ±637.56
WFC	562.95 ±173.07	-780.57 ±1118.2	389.75 ±737.45	NaN**	-478.31 ±2011.8
V	<u>353.09 ±117.31</u>	23.7 ±222.98	329.12 ±99.36	406.14 ±165.45	241.9 ±1070.96
XOM	<u>389.74 ±266.41</u>	356.79 ±252.04	-2.86 ±164.5	103.69 ±1097.92	2131.95 ±1424.47
PFE	-411.46 ±677.54	114.38 ±2267.64	<u>192.45 ±1656.14</u>	-163.88 ±1545.7	419.19 ±211.39
AMZN	121.02 ±244.41	<u>386.65 ±1305.54</u>	-597.9 ±1003.48	2047.26 ±2090.29	-1184.76 ±3526.17
BA	572.82 ±2258.07	16.24 ±853.6	-1211.11 ±658.4	-653.36 ±1339.05	143.52 ±725.58
META	1344.86 ± NaN*	-589.62 ±1253.38	1820.84 ±1903.45	-1695.45 ±3261.33	325.51 ±1269.18
MA	<u>122.96 ±65.76</u>	-99.61 ±266.94	246.72 ±69.7	-247.12 ±954.71	-577.71 ±893.93
CVS	1092.8 ±528.37	-634.47 ±774.1	989.94 ±761.13	795.5 ±367.36	-213.04 ±1318.44

Query Stock	Maximum Drawdown (%)				
	SimStock	TS2VEC	Corr1	Corr2	Coint
AAPL	-1.91 ±0.67	NaN**	-6.06 ±3.47	<u>-3.4 ±2.01</u>	NaN**
CMG	-2.99 ±2.54	-6.38 ±2.22	-14.98 ±6.37	-15.55 ±10.38	-20.05 ±15.48
MSFT	-2.82 ±1.88	<u>-5.67 ±0.24</u>	-7.16 ±4.7	-12.92 ±12.56	-13.57 ±2.58
WFC	-2.2 ±3.78	-9.2 ±8.24	<u>-5.76 ±4.54</u>	NaN**	-16.07 ±13.41
V	<u>-2.22 ±2.55</u>	-2.39 ±2.78	-0.43 ±0.75	-3.17 ±4.48	-7.48 ±5.89
XOM	<u>-3.4 ±0.41</u>	-3.03 ±2.97	-4.72 ±1.43	-6.55 ±6.9	-5.78 ±4.3
PFE	-7.84 ±7.41	-10.03 ±13.5	-9.99 ±9.42	-10.45 ±5.08	<u>-9.05 ±5.56</u>
AMZN	-5.06 ±4.99	-4.68 ±2.33	-15.14 ±9.38	-8.97 ±4.15	-31.92 ±13.14
BA	<u>-12.1 ±11.58</u>	-5.86 ±2.15	-15.48 ±8.63	-12.87 ±3.26	-14.44 ±7.33
META	-4.92 ±NaN*	<u>-8.77 ±8.96</u>	-11.59 ±4.95	-34.15 ±26.03	-25.43 ±11.24
MA	-2.18 ±2.51	-5.01 ±5.2	<u>-2.58 ±2.72</u>	-6.85 ±7.24	-11.51 ±4.68
CVS	<u>-3.02 ±1.88</u>	-9.2 ±8.57	<u>-4.57 ±4.6</u>	-2.92 ±1.36	-19.87 ±17.05

Table 3. Average terminal wealth (first row) and maximum drawdown (MDD) (second row) achieved by applying pairs trading to the top@3 similar stocks identified by SimStock, TS2VEC, Corr1, Corr2, coint for each query stock. The best performing method is highlighted in bold, while the second best is in underline.

Note: NaN** values in both the terminal wealth and MDD indicate that the method failed to generate buy/sell signals for all three stocks in the pair. NaN* values only in the standard deviation indicate that the method failed to generate buy/sell signals for two out of the three stocks in the pair. For all other values, all method generated buy/sell signals for all three stocks in the pair.

4.4. Application to index tracking of thematic ETFs

Thematic ETFs have gained popularity among retail investors who wish to invest in specific themes or trends, such as innovative technologies. These ETFs are composed of assets that share a common theme, providing a contrast to traditional sector-based classifications, which can be slow and ambiguous in reflecting technological advancements and innovations. As thematic ETFs aim to identify innovative tech companies, many investors are interested in tracking these ETFs.

In this section, we compare the performance of `SimStock` and other baselines in identifying stocks to track four popular thematic ETFs. Our goal is to find similar stocks from the US, SSE, SZSE, and TSE exchanges that can effectively track the performance of these thematic ETFs. In other words, we use the thematic ETFs as queries to find similar stocks in each model.

For this analysis, we selected the four popular thematic ETFs: ARK Innovation ETF (ARKK), First Trust Cloud Computing ETF (SKYY), Global X Robotics & AI ETF (BOTZ), and Global X Lithium & Battery Tech ETF (LIT). The in-sample and out-of-sample periods are the same as those used in section 4.1.

4.4.1. Evaluation metrics. To evaluate the performance of `SimStock` and the baselines in tracking thematic ETFs, we first find the top k similar stocks based on the thematic ETF queries. We then create equal-weighted portfolios using these similar stocks to track the thematic ETFs. We use two metrics to assess the tracking performance: Tracking Error (TE) and Tracking Error Volatility (TEV) (Focardi and Fabozzi 2004).

The end-of-the-period tracking error (TE) is calculated as follows:

$$\text{TE} = \sqrt{\frac{1}{n} \sum_{j=1}^n (R_j^I - R_j^P)^2} \quad (13)$$

where R_j^I and R_j^P are the cumulative returns of the query ETF and the tracking portfolio at period j , respectively, and n is the number of periods.

Tracking Error Volatility (TEV) measures the volatility of the difference in returns between the tracking portfolio and the target ETF. It is calculated as the standard deviation of the difference in returns:

$$\text{TEV} = \sqrt{\frac{1}{n} \sum_{j=1}^n (R_j^I - R_j^P - \mathbb{E}[R^I - R^P])^2} \quad (14)$$

where $\mathbb{E}[R^I - R^P]$ is the mean difference in returns between the target ETF and the tracking portfolio over the n periods. For each thematic ETF, we find the top@ k similar stocks on each exchange (US, SSE, SZSE, TSE) using `SimStock` and the baseline methods. We then create equal-weighted portfolios using these similar stocks to track the performance of the thematic ETFs.

4.4.2. Quantitative evaluation. In this subsection, we quantitatively evaluate the performance of `SimStock` and the baseline methods in tracking the four thematic ETFs. Tables 4 report the tracking error (TE) for each ETF using the top@ k similar stocks identified by each method across the four exchanges. The best performing method for each k and exchange is highlighted in bold, while the second best is in underline. The results of the evaluation with Tracking Error Volatility (TEV) can be found in Appendix B.

A clear pattern emerges from the results: `SimStock` consistently achieves the bold (blue) or

second lowest (underline) tracking error for the majority of k values across three of the four ETFs (ARKK, SKYY, and BOTZ) and exchanges. This strong performance demonstrates that the similar stocks identified by `SimStock` are highly effective at tracking these thematic ETFs. In contrast to `SimStock`, while the `Corr1` and `Corr2` baselines show reasonable performance in some cases, they lack consistency across different scenarios. `TS2VEC`, on the other hand, generally exhibits higher tracking errors, indicating that the stocks it identifies do not track the ETFs as closely.

Notably, `SimStock`'s performance on the US exchange is particularly impressive for ARKK, SKYY, and BOTZ. It substantially outperforms the other methods, especially for lower values of k (e.g., top 10-20 similar stocks), suggesting that `SimStock` excels at identifying a concentrated set of the most relevant stocks for tracking these ETF themes. `SimStock`'s ability to identify highly relevant stocks in this market, even with a smaller number of stocks (low top@ k values), is a testament to its effectiveness and precision in capturing the key components driving these thematic ETFs. While the performance gap between `SimStock` and the other methods is less pronounced on the other exchanges, `SimStock` still achieves the best or second-best performance in most cases for ARKK, SKYY, and BOTZ. However, `SimStock`'s performance in tracking the Global X Lithium & Battery Tech ETF (LIT) is relatively weaker compared to the other ETFs.

In addition to the quantitative evaluation using tracking error and tracking error volatility, we can also visually assess the performance of `SimStock` and the baseline methods in tracking the thematic ETFs by examining their cumulative return curves. Figure 5 showcases the cumulative return curves for each thematic ETF (ARKK, SKYY, BOTZ, and LIT) and the corresponding tracking portfolios created using the top@10 similar stocks identified by each method on the US exchange.

The cumulative return curves provide a clear visual representation of how closely the tracking portfolios match the performance of their respective thematic ETFs over time. For ARKK, SKYY, and BOTZ, the `SimStock` portfolio (shown in blue) closely follows the trajectory of the thematic ETF (shown in dotted black), indicating that the top 10 stocks identified by `SimStock` effectively capture the performance of these ETFs. In contrast, the portfolios based on the other methods (`TS2VEC` in purple, `Corr1` in green, and `Corr2` in red) exhibit larger deviations from the ETF curves, suggesting that their selected stocks are less effective at tracking the ETFs.

The cumulative return curves represent `SimStock`'s ability to identify a concentrated set of stocks that effectively track the performance of thematic ETFs, particularly for ARKK, SKYY, and BOTZ, based on the top@10 similar stocks from the US exchange.

ARK Innovation ETF (ARKK)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.2828	0.1363	0.0746	0.0822	0.0649	0.0592
	Corr2	0.4264	0.2431	0.1132	0.1455	0.0963	0.1149
	TS2VEC	0.4752	0.5772	0.5321	0.4887	0.3723	0.3046
	SimStock	0.1934	0.0764	0.0975	0.1785	0.1822	0.2414
SSE	Corr1	0.4014	0.4244	0.4191	0.4250	0.4333	0.4365
	Corr2	0.2753	0.2651	0.2739	0.2883	0.2681	0.2747
	TS2VEC	0.3813	0.3815	0.3748	0.4037	0.4030	0.4142
	SimStock	0.3064	0.3133	0.2778	0.2774	0.2754	0.3069
SZSE	Corr1	0.3745	0.3843	0.3607	0.4006	0.3843	0.3923
	Corr2	0.2302	0.2063	0.2138	0.2208	0.2380	0.2547
	TS2VEC	0.2793	0.2843	0.3132	0.3074	0.2993	0.3010
	SimStock	0.1621	0.1643	0.1959	0.2252	0.2354	0.2272
TSE	Corr1	0.3965	0.3214	0.2694	0.2533	0.2591	0.2573
	Corr2	0.1458	0.1707	0.1850	0.1846	0.1869	0.2002
	TS2VEC	0.2712	0.2940	0.2574	0.2856	0.2841	0.2790
	SimStock	0.1440	0.1322	0.1310	0.1328	0.1262	0.1321

First Trust Cloud Computing ETF (SKYY)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.1502	0.1968	0.2064	0.1937	0.1310	0.1235
	Corr2	0.3643	0.2941	0.2130	0.2137	0.1673	0.1482
	TS2VEC	0.2525	0.2370	0.2104	0.2377	0.2262	0.2263
	SimStock	0.0939	0.0366	0.0645	0.0814	0.0603	0.0596
SSE	Corr1	0.2990	0.3471	0.3612	0.3715	0.3757	0.3872
	Corr2	0.1865	0.1961	0.1941	0.1554	0.1647	0.1760
	TS2VEC	0.2939	0.3475	0.3857	0.3616	0.3572	0.3240
	SimStock	0.1607	0.1291	0.1688	0.1615	0.1748	0.1912
SZSE	Corr1	0.4157	0.4014	0.3893	0.3915	0.3771	0.3603
	Corr2	0.1329	0.1390	0.1532	0.1617	0.1679	0.1771
	TS2VEC	0.2388	0.2789	0.2616	0.2598	0.2690	0.2499
	SimStock	0.1230	0.1543	0.1731	0.1651	0.1559	0.1455
TSE	Corr1	0.1363	0.1822	0.1756	0.2212	0.2206	0.2300
	Corr2	0.1118	0.1489	0.1232	0.1389	0.1321	0.1261
	TS2VEC	0.2220	0.2395	0.2384	0.2276	0.2290	0.2369
	SimStock	0.0940	0.0733	0.0734	0.0660	0.0593	0.0691

Global X Robotics & AI ETF (BOTZ)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.2847	0.1581	0.0981	0.0881	0.0784	0.0897
	Corr2	0.0946	0.1526	0.1249	0.1216	0.1448	0.1373
	TS2VEC	0.1890	0.3221	0.3368	0.2763	0.2833	0.3247
	SimStock	0.0364	0.0647	0.0946	0.0910	0.0890	0.0818
SSE	Corr1	0.2912	0.2870	0.2831	0.2738	0.2745	0.2815
	Corr2	0.0863	0.1205	0.1198	0.1179	0.1220	0.1382
	TS2VEC	0.2512	0.2535	0.2638	0.2955	0.2646	0.2696
	SimStock	0.1471	0.1520	0.1514	0.1573	0.1509	0.1283
SZSE	Corr1	0.2259	0.2063	0.2329	0.2466	0.2645	0.2763
	Corr2	0.1016	0.1142	0.1143	0.1144	0.1242	0.1169
	TS2VEC	0.1842	0.1516	0.1527	0.1387	0.1376	0.1694
	SimStock	0.0812	0.1289	0.1171	0.1246	0.1331	0.1366
TSE	Corr1	0.1646	0.1420	0.1457	0.1454	0.1411	0.1236
	Corr2	0.0951	0.0916	0.0863	0.0844	0.0910	0.0895
	TS2VEC	0.2155	0.2343	0.1943	0.1952	0.1959	0.1843
	SimStock	0.0639	0.0730	0.0712	0.0797	0.0802	0.0764

Global X Lithium & Battery Tech ETF (LIT)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.2167	0.2819	0.2512	0.2638	0.2732	0.2465
	Corr2	0.1424	0.2337	0.2234	0.2195	0.2512	0.1669
	TS2VEC	0.0838	0.1094	0.0859	0.0924	0.1080	0.1070
	SimStock	0.0446	0.0557	0.0647	0.0666	0.0694	0.0703
SSE	Corr1	0.1756	0.1671	0.1485	0.1431	0.1354	0.1294
	Corr2	0.1159	0.1061	0.0983	0.0923	0.0899	0.0905
	TS2VEC	0.0715	0.0798	0.0772	0.0790	0.0817	0.0803
	SimStock	0.1003	0.1355	0.1084	0.1009	0.1161	0.1211
SZSE	Corr1	0.1562	0.1431	0.1676	0.1676	0.1700	0.1730
	Corr2	0.2638	0.2566	0.2073	0.1906	0.1838	0.1690
	TS2VEC	0.1504	0.1236	0.1138	0.1090	0.0997	0.0962
	SimStock	0.2861	0.3187	0.2833	0.2892	0.2734	0.2412
TSE	Corr1	0.1839	0.2430	0.2477	0.2769	0.2600	0.2242
	Corr2	0.3249	0.3030	0.2580	0.2457	0.2826	0.2662
	TS2VEC	0.1563	0.1537	0.1642	0.1727	0.1502	0.1768
	SimStock	0.2557	0.2794	0.2763	0.2866	0.2861	0.2735

Table 4. Tracking errors (TE) of SimStock and baseline methods for tracking the performance of four thematic ETFs (ARKK, SKYY, BOTZ, and LIT) using top k similar stocks from the US, SSE, SZSE, and TSE exchanges. The best performing method for each k and exchange is highlighted in bold, while the second best is in underline.




Figure 5. Cumulative return curves of the four thematic ETFs (ARKK, SKYY, BOTZ, and LIT) and their corresponding tracking portfolios constructed using the top 10 similar stocks identified by `SimStock` and the baseline methods (TS2VEC, Corr1, and Corr2) from the US exchange. The closer a portfolio’s curve follows the respective ETF curve (dotted black line), the better the tracking performance.

4.4.3. Qualitative evaluation. For the ARK Innovation ETF (ARKK), which focuses on disruptive innovation, `SimStock` identifies stocks of companies involved in cutting-edge technologies such as 3D computing (MLTS, DDD), fintech company (SQ), and artificial intelligence (SHOP). Similarly, for the First Trust Cloud Computing ETF (SKYY), `SimStock` finds stocks of leading cloud computing (NET, DBX, TYL, GWRE and GLOB) and software-as-a-service providers (INTU, ESTC and PAYC). The qualitative assessment of `SimStock`’s similar stock selections for all thematic ETFs can be found in Appendix B.

4.5. Application to Portfolio optimization

In this section, we investigate whether `SimStock` embeddings can enhance portfolio optimization. Specifically, we construct the correlation matrix using the `SimStock` embedding as a similarity measure, and use it as an input for portfolio optimization. We compare the portfolio performance using the `SimStock` embedding with other covariance estimators. Notice that `SimStock` is designed to identify stocks with similar behavior, but it does not actively search for stocks that exhibit completely different dynamics compared to a given stock. This may limit its ability to ensure optimal portfolio diversification.

4.5.1. Optimization formulation. We consider the Mean-Variance Optimization (MVO) framework (Markowitz 1952) for portfolio optimization. The MVO approach determines the optimal asset weights (ω_i) in a portfolio, subject to constraints on risk and return, assuming the expected

return (μ_i), variance (σ_{ii}^2), and covariance (σ_{ij}) of each asset are known. This study adopts the experimental design and methodology established by (Gerber *et al.* 2021) We randomly sampled the asset classes 100 times with 10, 30 and 50 stocks each from the S&P500 and JPX Prime 150.

The long-only MVO problem, accounting for transaction costs, is formulated as:

$$\begin{aligned} & \text{Maximize: } w^T \mu - \psi 1^T |w - w_0| \\ & \text{Subject to: } w^T \Sigma w \leq \sigma_{\text{target}}^2, \quad w^T 1 = 0, \quad 0 \leq w_k \leq 1, \forall k = 1, 2, \dots, N. \end{aligned} \quad (15)$$

The objective function seeks to maximize the portfolio's expected return ($\omega^T \mu$) while minimizing transaction costs, represented by the term $\psi 1^T |\omega - \omega_0|$. Here, ψ denotes the fixed proportional transaction cost, set to 10 basis points (0.1%) in this study, and ω_0 represents the previous portfolio weights. By including this term, we penalize deviations from the previous weights, effectively controlling portfolio turnover and reducing trading costs.

The optimization is subject to three constraints: (1) the portfolio variance must not exceed a predetermined risk target (σ_{target}^2), (2) the sum of portfolio weights must equal unity, and (3) individual asset weights must be non-negative and not exceed one, prohibiting short selling. By solving this constrained optimization problem, we obtain the optimal portfolio weights ω^* for a given risk target σ_{target} and transaction cost ψ . The set of optimal portfolios across various risk levels constitutes the efficient frontier, with each point representing the highest achievable return for a specific level of risk.

We estimate the expected return μ_{ti} for asset i at time t using the sample mean of its historical returns over a T -month lookback window. For the covariance matrix Σ , we compare the performance of historical covariance (HC), shrinkage method (SM) (Ledoit and Wolf 2004a), Gerber statistic (GS) (Gerber *et al.* 2021), TS(TS2VEC) (Yue *et al.* 2022) embedding and SimStock embedding (SS).

4.5.2. Portfolio backtesting procedure. To evaluate the performance of the different covariance estimators in the context of portfolio optimization, we employ the following backtesting procedure. Starting from Jan 2021, at the beginning of each month, we use the monthly returns of the current asset universe over the previous $T = 12$ months to estimate the expected return vector μ and the covariance matrix Σ .

We then apply a quadratic optimizer to solve for the optimal portfolio weights ω^* given a specific risk target σ_{target} for MVO and the optimal portfolio weights w^* that minimize the portfolio variance for MVP. The previous portfolio is rebalanced according to the new optimal weights, and this optimized portfolio is held for one month. At the end of the month, the realized portfolio return is calculated as $w^{*T} \tilde{r}$, where \tilde{r} represents the vector of realized asset returns for that month.

This process is repeated by rolling the in-sample period forward by one month and computing the updated efficient portfolio for the next month. The rolling-window approach allows the portfolios to adapt to structural changes in the market and mitigates data mining bias. As the initialization of the first portfolio requires two years of monthly returns, our performance evaluation covers the period from Jan 2022 to Feb 2024.

4.5.3. Evaluation for portfolio optimization. In this subsection, to evaluate the performance of portfolios optimized using the SimStock embedding covariance matrix (Σ_{SS}), we compare their ex-post efficient frontiers with those of portfolios constructed using traditional covariance estimators, namely the historical covariance (HC), shrinkage method (SM), Gerber statistic (GS) and TS2VEC (TS), denoted as Σ_{HC} , Σ_{SM} , Σ_{GS} and Σ_{TS} , respectively. To obtain the covariance matrix Σ_{SS} using the TS2VEC and SimStock embedding, we first compute the pairwise L2 distances between the embeddings of each stock pair. These distances are then scaled to the range $[-1, 1]$ using

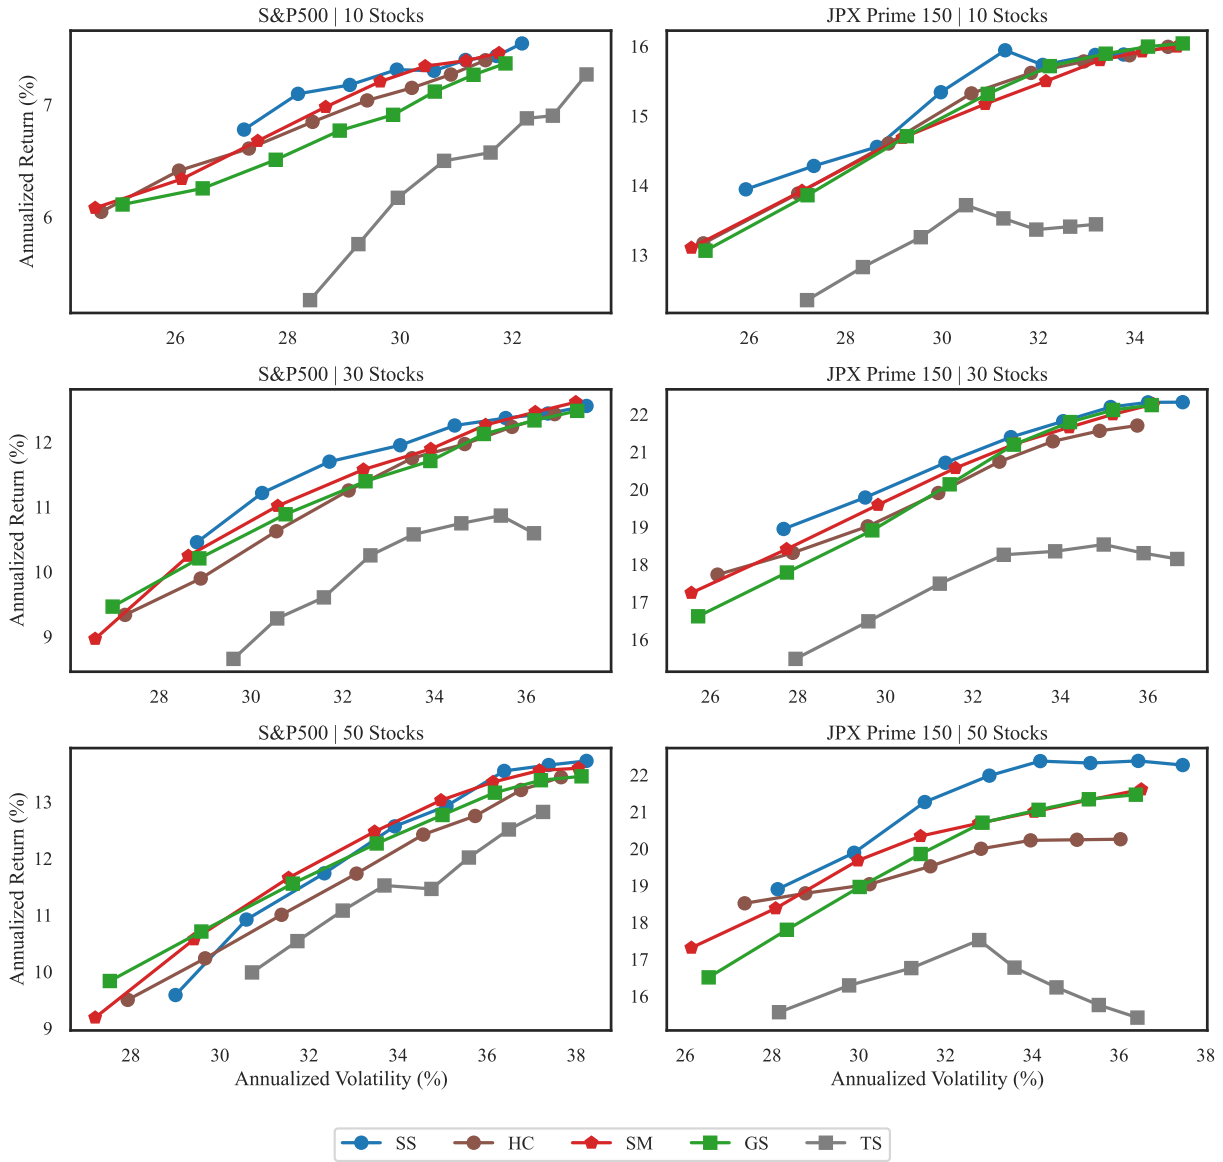


Figure 6. Ex-post efficient frontiers displaying annualized return and volatility of portfolios optimized for different risk targets. The black vertical dotted lines represent the average volatility of the S&P500 and JPX Prime 150, respectively.

max-min normalization, which brings the stocks closer to 1 in the correlation matrix for smaller L2 distances, indicating greater similarity between stocks. The resulting normalized distances are used to populate the correlation matrix, represents the correlation coefficient between stocks i and j as estimated by embeddings. We then compute the covariance matrix Σ_{TS} and Σ_{SS} . The analysis is conducted on portfolios of varying sizes (10, 30, and 50 stocks) from both the S&P500 and JPX Prime 150 universes, across different risk targets.

As shown in Figure 6, the **SimStock** embedding approach slightly outperforms the other covariance estimators across all portfolio sizes and risk levels in both the S&P500 and JPX Prime 150 universes. The portfolios constructed using Σ_{SS} achieve higher annualized returns for a given level of annualized volatility compared to those based on Σ_{HC} , Σ_{SM} , Σ_{GS} and Σ_{TS} . This outperformance is slightly more evident for smaller portfolio sizes (10 and 30 stocks), where the **SimStock** embedding

method demonstrates an advantage in capturing meaningful relationships between assets, leading to superior risk-adjusted performance. While our approach did not aim to take into account the correlation coefficients between all stocks, these results suggest that `SimStock` has the potential to slightly improve portfolio optimization results in portfolio construction settings.

The underperformance of `TS2VEC` is particularly noticeable when compared to the `SimStock` (SS) embedding approach, which consistently outperforms all other methods. The gap between the TS and SS curves is most evident for the smaller portfolio sizes (10 and 30 stocks), suggesting that the `TS2VEC` embeddings may struggle to capture the nuanced relationships between assets in these more concentrated portfolios.

4.5.4. Performance metrics for portfolio optimization. This section compares the performance of five portfolio construction methods (`Simstock` embedding, Historical Covariance, Shrinkage Method, Gerber Statistic, and `TS2VEC` embedding) for MVO, using 30-stock portfolios from the S&P 500 and JPX Prime 150 universes. See Appendix C for results for 10 and 50 stocks.

Tables 5 and 6 present the performance metrics for the target risk optimizations in the S&P500 and JPX Prime 150 universes, respectively, at four different risk target levels (24%, 27%, 30%, and 33%). See Appendix C for results for 10 and 50 stocks.

In the S&P500 universe (Table 5), the `Simstock` embedding (SS) method produced the highest returns at all risk target levels while maintaining risk measures comparable to the other methods. Returns and risk measures increased for all methods as the risk target rose. The Historical Covariance (HC), Shrinkage Method (SM), and Gerber Statistic (GS) methods performed similarly to each other, but slightly worse than SS, particularly at higher risk levels. For example, at the 24% risk target, SS achieved an 11.31% arithmetic return compared to 10.04%, 9.61%, and 10.14% for HC, SM, and GS respectively. At the 33% risk target, SS's return increased to 13.33%, while HC, SM and GS saw returns of 12.50%, 12.85% and 12.69%. The `TS2VEC` embedding (TS) method consistently underperformed the others, with a 9.40% return at the 24% risk level and 11.49% at the 33% level, below all other methods.

The JPX Prime 150 universe (Table 6) showed similar patterns, with SS producing the best returns, followed by HC, GS, SM and TS in that order across risk levels. SS achieved a 20.12% return at the 24% risk target and 22.47% at the 33% target. In contrast, TS had the lowest returns at 14.60% and 18.45% for those risk levels. The HC, GS, and SM methods produced returns that were higher than TS but lower than SS at each corresponding risk level.

Despite potentially higher volatility which is evident in the ex-post efficient frontiers in Figure 6, the performance of the SS method highlights its potential as a meaningful and effective stock embedding technique for portfolio optimization.

S&P500 — 30 Stocks	Target Volatility (24%)				Target Volatility (27%)				Target Volatility (30%)				Target Volatility (33%)			
	SS	HC	SM	TS	SS	HC	SM	TS	SS	HC	SM	TS	SS	HC	SM	TS
Covariance Method	11.31	10.04	9.61	10.14	12.22	10.75	11.02	11.05	10.51	12.84	11.66	12.03	11.97	13.33	12.50	12.85
Arithmetic Return (%)	10.39	9.30	8.97	9.43	11.158	9.86	10.21	10.17	9.43	11.60	10.57	10.967	10.83	11.88	11.18	11.32
Geometric Return (%)	37.36	33.07	31.60	33.40	40.54	35.40	36.63	36.49	31.04	42.64	38.44	39.930	39.61	44.18	41.27	42.05
Cumulative Return (%)	28.82	27.27	26.61	27.00	30.24	28.92	28.66	28.90	30.24	31.72	30.58	30.623	30.79	33.25	32.17	32.49
Annualized SD (%)	-0.12	-0.15	-0.12	-0.122	-0.14	-0.17	-0.16	-0.16	-0.21	-0.18	-0.20	-0.191	-0.20	-0.22	-0.22	-0.23
Annualized Skewness	3.17	3.22	3.23	3.161	3.22	3.28	3.29	3.24	2.86	3.27	3.33	3.343	3.28	3.27	3.33	3.28
Annualized Kurtosis	-24.90	-23.96	-23.47	-23.59	-25.65	-25.36	-24.57	-25.04	-25.93	-26.69	-26.44	-25.971	-26.57	-28.20	-27.63	-27.57
Maximum Drawdown (%)	-10.44	-10.22	-9.99	-10.05	-10.77	-10.63	-10.53	-10.52	-11.38	-11.19	-10.95	-10.921	-11.01	-11.72	-11.40	-11.48
Monthly 95% VaR (%)	0.44	0.40	0.39	0.42	0.46	0.41	0.43	0.43	0.31	0.46	0.42	0.447	0.43	0.45	0.43	0.36
Sharpe Ratio	8.68	8.39	8.49	8.36	8.69	8.48	8.56	8.49	8.04	8.73	8.54	8.592	8.56	8.67	8.57	8.54
Annualized Turnover																

Table 5. This table presents the performance metrics for four portfolio construction methods in the S&P500: Simstock embedding (SS), historical covariance (HC), shrinkage method (SM), Gerber statistic (GS) and TS2VEC embedding (TS). The portfolios were optimized for four different risk target levels: 24%, 27%, 30%, and 33%. The performance was evaluated over the full testing period from January 2022 to February 2024. The 3-month U.S. Treasury Bill rate was used as the risk-free rate for performance calculations. Transaction costs were modeled as 10 basis points of the traded volume for each rebalancing event.

JPX Prime 150 — 30 Stocks	Target Volatility (24%)				Target Volatility (27%)				Target Volatility (30%)				Target Volatility (33%)			
	SS	HC	SM	TS	SS	HC	SM	TS	SS	HC	SM	TS	SS	HC	SM	TS
Covariance Method	20.12	19.02	18.61	17.80	20.85	19.78	20.02	19.20	15.67	21.70	20.65	21.40	20.51	22.47	21.70	22.50
Arithmetic Return (%)	18.63	17.72	17.24	16.62	19.24	18.31	18.40	17.78	14.25	19.96	19.01	19.58	18.90	20.62	19.90	20.57
Geometric Return (%)	71.15	66.53	64.60	61.55	74.32	69.56	70.25	67.23	49.12	77.95	73.23	76.27	72.90	81.67	77.87	81.37
Cumulative Return (%)	26.83	26.16	25.56	25.71	28.51	27.88	27.74	27.74	28.39	29.99	29.59	29.83	29.69	31.33	31.20	31.59
Annualized SD (%)	0.28	0.16	0.17	0.13	0.32	0.17	0.21	0.18	0.15	0.34	0.19	0.23	0.21	0.32	0.19	0.23
Annualized Skewness	3.37	2.94	2.99	2.92	3.43	3.02	3.08	3.01	2.67	3.49	3.06	3.17	3.09	3.51	3.12	3.16
Annualized Kurtosis	-19.17	-19.63	-19.52	-19.44	-20.37	-20.89	-20.85	-20.62	-22.89	-21.16	-22.09	-22.15	-21.96	-22.57	-23.28	-23.29
Maximum Drawdown (%)	-8.52	-8.89	-8.49	-8.82	-9.00	-9.46	-9.20	-9.36	-9.92	-9.41	-9.95	-9.79	-9.90	-9.79	-10.46	-10.31
Monthly 95% VaR (%)	0.96	0.92	0.91	0.86	0.93	0.89	0.90	0.86	0.65	0.92	0.87	0.90	0.87	0.92	0.87	0.88
Sharpe Ratio	8.81	8.38	8.50	8.54	8.83	8.52	8.57	8.62	8.15	8.81	8.59	8.56	8.59	8.85	8.58	8.54
Annualized Turnover																

Table 6. This table presents the performance metrics for four portfolio construction methods in the JPX Prime 150: Simstock embedding (SS), historical covariance (HC), shrinkage method (SM), Gerber statistic (GS) and TS2VEC embedding (TS). The portfolios were optimized for four different risk target levels: 24%, 27%, 30%, and 33%. The performance was evaluated over the full testing period from January 2022 to February 2024. The 3-month U.S. Treasury Bill rate was used as the risk-free rate for performance calculations. Transaction costs were modeled as 10 basis points of the traded volume for each rebalancing event.

4.5.5. Characteristics of SimStock similarity matrix. To investigate what distinguishes the SimStock embedding method from other portfolio optimization methods, we can compare the correlation matrices generated by different methodologies. Specifically, we propose the following mapping:

$$\frac{\| \text{MD} - \text{RC}^{\text{future}} \|_{\text{F}}}{\| \text{MD} - \text{RC}^{\text{past}} \|_{\text{F}}} \leq 1 \quad (16)$$

Here, MD refers to the correlation matrix obtained using a specific methodology (e.g., SS, SM, GS or TS), while $\text{RC}^{\text{future}}$ and RC^{past} represent the realized correlation matrices for the future and past periods, respectively.

If the value of this mapping is less than or equal to 1, it implies that $\| \text{MD} - \text{RC}^{\text{future}} \|_{\text{F}} \leq \| \text{MD} - \text{RC}^{\text{past}} \|_{\text{F}}$. This means that the current methodology is tracking the future correlation matrix well, as the Frobenius distance between the method's correlation matrix and the future realized correlation matrix is smaller than the distance between the method's correlation matrix and the past correlation matrix.




Figure 7. Comparison of Frobenius distances between correlation matrices generated by different methodologies (SS: SimStock embedding, SM: shrinkage method, GS: Gerber statistic) and realized correlation matrices for consecutive years in the S&P500 and JPX Prime 150 universes. Lower values indicate better tracking of the future correlation matrix. The SS method consistently exhibits lower Frobenius distances compared to the SM, GS and TS methods across most month pairs in both universes, suggesting its superior ability to capture future asset correlations.

Figure 7 illustrates the Frobenius distances for different methodologies in the S&P500 and JPX Prime 150 universes. The figure compares the correlation matrices between the same months in consecutive years (e.g., Feb-2021 vs Feb-2022, Mar-2021 vs Mar-2022, etc.). For example, "Feb-2021 vs Feb-2022" means that the correlation series for the period from February 2021 to January 2022 is compared to the correlation series for the period from February 2022 to January 2023, representing the current year and the future year, respectively.

For the S&P500, the Frobenius distances for the SimStock embedding method (SS) are consistently lower than those of the shrinkage method (SM), Gerber statistic (GS) and TS2VEC embedding (TS) across all month pairs. This suggests that the SS method is better at tracking the future correlation matrix compared to the SM and GS methods.

Similarly, in the JPX Prime 150 universe, the SS method generally exhibits lower Frobenius distances than the SM, GS and TS methods for most month pairs. This indicates that the SS method is more effective at capturing the future correlation structure of the assets compared to the other methodologies.

The ability of the SimStock embedding method to better track the future correlation matrix can be attributed to its unique approach, temporal representation learning, in capturing the latent relationships between assets. This improved understanding of the underlying asset dynamics could enable the SS method to generate correlation matrices that more closely resemble the future realized correlations, thereby providing an advantage edge in portfolio optimization.

5. Conclusion

In this paper, we propose **SimStock**, a novel temporal self-supervised learning framework that aims to learn robust and informative representations of financial time series data. By incorporating techniques from SSL and temporal domain generalization, **SimStock** captures the complex relationships between different financial assets while accounting for temporal shifts in the data distribution. The proposed dimension corruption method integrates temporal patterns into the corruption process, enabling **SimStock** to learn representations that are robust to noise and non-stationarity in the data.

We conduct extensive experiments on four real-world benchmarks with thousands of stocks to demonstrate the effectiveness of **SimStock** in finding similar stocks. Our results show that **SimStock** consistently outperforms existing methods in both one-to-one and one-to-many scenarios, achieving state-of-the-art performance in terms of accuracy and robustness. The qualitative analysis further highlights the ability of **SimStock** to identify stocks with similar fundamental characteristics and industry-specific similarities, simplifying the process of screening potential investment opportunities.

The practical utility of **SimStock** is demonstrated through various financial applications, including pairs trading, index tracking, and portfolio optimization. In the pairs trading experiment, using the similar stocks identified by **SimStock** leads to superior profitability compared to traditional methods. For index tracking, **SimStock** exhibits lower tracking errors and better alignment with the underlying themes of thematic ETFs, particularly for ARKK, SKYY, and BOTZ. Finally, in portfolio optimization, the **SimStock** embedding approach slightly outperforms traditional covariance estimators, achieving higher risk-adjusted returns across different portfolio sizes and risk levels.

The success of **SimStock** can be attributed to its unique approach of temporal representation learning, which captures the latent relationships between assets and better tracks future correlation matrices compared to other methodologies. This improved understanding of the underlying asset dynamics provides an edge in various financial applications.

References

- Ang, A. and Bekaert, G., International asset allocation with regime shifts. *The review of financial studies*, 2002, **15**, 1137–1187.
- Bai, G., Ling, C. and Zhao, L., Temporal Domain Generalization with Drift-Aware Dynamic Neural Networks. *arXiv preprint arXiv:2205.10664*, 2022.
- Balntas, V., Riba, E., Ponsa, D. and Mikolajczyk, K., Learning local feature descriptors with triplets and shallow convolutional neural networks.. In *Proceedings of the Bmvc*, Vol. 1, p. 3, 2016.
- Barberis, N., Investing for the long run when returns are predictable. *The Journal of Finance*, 2000, **55**, 225–264.
- Berndt, D.J. and Clifford, J., Using dynamic time warping to find patterns in time series.. In *Proceedings of the KDD workshop*, Vol. 10, pp. 359–370, 1994.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P. and Joulin, A., Emerging properties in self-supervised vision transformers. In *Proceedings of the Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- Chen, T., Kornblith, S., Norouzi, M. and Hinton, G., A simple framework for contrastive learning of visual

- representations. In *Proceedings of the International conference on machine learning*, pp. 1597–1607, 2020.
- Choi, H. and Kang, P., Multi-task self-supervised time-series representation learning. *arXiv preprint arXiv:2303.01034*, 2023.
- Chopra, V.K. and Ziemba, W.T., The effect of errors in means, variances, and covariances on optimal portfolio choice. In *Handbook of the fundamentals of financial decision making: Part I*, pp. 365–373, 2013, World Scientific.
- Chung, M., Lee, Y., Kim, J.H., Kim, W.C. and Fabozzi, F.J., The effects of errors in means, variances, and correlations on the mean-variance framework. *Quantitative Finance*, 2022, **22**, 1893–1903.
- Cont, R., Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 2001, **1**, 223.
- Devlin, J., Chang, M.W., Lee, K. and Toutanova, K., Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Enders, W., Cointegration and error-correction models. *Applied econometrics time series*, 2004, pp. 319–386.
- Fama, E.F. and French, K.R., Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 1993, **33**, 3–56.
- Fama, E.F. and French, K.R., A five-factor asset pricing model. *Journal of financial economics*, 2015, **116**, 1–22.
- Fang, Y., Dong, L., Bao, H., Wang, X. and Wei, F., Corrupted image modeling for self-supervised visual pre-training. *arXiv preprint arXiv:2202.03382*, 2022.
- Focardi, S.M. and Fabozzi, F.J., A methodology for index tracking based on time-series clustering. *Quantitative Finance*, 2004, **4**, 417–425.
- Gatev, E., Goetzmann, W.N. and Rouwenhorst, K.G., Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, 2006, **19**, 797–827.
- Gerber, S., Markowitz, H., Ernst, P., Miao, Y., Sargen, P. *et al.*, The Gerber statistic: A robust co-movement measure for portfolio optimization. *No and Sargen, Paul, The Gerber Statistic: A Robust Co-Movement Measure for Portfolio Optimization (July 4, 2021)*, 2021.
- Gorishniy, Y., Rubachev, I., Khrulkov, V. and Babenko, A., Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems*, 2021, **34**, 18932–18943.
- Grill, J.B., Strub, F., Althé, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M. *et al.*, Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 2020, **33**, 21271–21284.
- Guidolin, M. and Timmermann, A., Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 2007, **31**, 3503–3544.
- He, K., Fan, H., Wu, Y., Xie, S. and Girshick, R., Momentum contrast for unsupervised visual representation learning. In *Proceedings of the Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.
- Hochreiter, S. and Schmidhuber, J., Long short-term memory. *Neural computation*, 1997, **9**, 1735–1780.
- Hwang, Y., Park, J., Lee, Y. and Lim, D.Y., Stop-loss adjusted labels for machine learning-based trading of risky assets. *Finance Research Letters*, 2023, **58**, 104285.
- Jenni, S., Meishvili, G. and Favaro, P., Video representation learning by recognizing temporal transformations. In *Proceedings of the Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, pp. 425–442, 2020.
- Jiang, A.Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D.S., Casas, D.d.l., Hanna, E.B., Bressand, F. *et al.*, Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- Jing, L. and Tian, Y., Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2020, **43**, 4037–4058.
- Jobson, J.D. and Korkie, B., Estimation for Markowitz efficient portfolios. *Journal of the American Statistical Association*, 1980, **75**, 544–554.
- Kan, R. and Zhou, G., Optimal estimation for economic gains: Portfolio choice with parameter uncertainty. *Washington University, St. Louis*, 2004.
- Kan, R. and Zhou, G., Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 2007, **42**, 621–656.
- Ledoit, O. and Wolf, M., Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of empirical finance*, 2003, **10**, 603–621.
- Ledoit, O. and Wolf, M., Honey, I Shrunk the Sample Covariance Matrix. *Journal of Portfolio Management*,

- 2004a, **30**, 110.
- Ledoit, O. and Wolf, M., A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 2004b, **88**, 365–411.
- Lee, K., Chang, M.W. and Toutanova, K., Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*, 2019.
- Lintner, J., The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. In *Stochastic optimization models in finance*, pp. 131–155, 1975, Elsevier.
- Lo, A., *Adaptive markets: Financial evolution at the speed of thought*, 2017, Princeton University Press.
- Lu, J., Liu, A., Dong, F., Gu, F., Gama, J. and Zhang, G., Learning under concept drift: A review. *IEEE transactions on knowledge and data engineering*, 2018, **31**, 2346–2363.
- Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S. *et al.*, Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Markowitz, H.M., Portfolio selection. *Journal of Finance*, 1952, **7**, 71–91.
- Markowitz, H., Portfolio Selection. *The Journal of Finance*, 1952, **7**, 77–91.
- Nelson, C.R. and Plosser, C.R., Trends and random walks in macroeconomic time series: some evidence and implications. *Journal of monetary economics*, 1982, **10**, 139–162.
- Pan, S.J. and Yang, Q., A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 2009, **22**, 1345–1359.
- Qian, R., Meng, T., Gong, B., Yang, M.H., Wang, H., Belongie, S. and Cui, Y., Spatiotemporal contrastive video representation learning. In *Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6964–6974, 2021.
- Qiu, X., Sun, T., Xu, Y., Shao, Y., Dai, N. and Huang, X., Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 2020, **63**, 1872–1897.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I. *et al.*, Language models are unsupervised multitask learners. *OpenAI blog*, 2019, **1**, 9.
- Rapach, D.E., Strauss, J.K. and Zhou, G., Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies*, 2010, **23**, 821–862.
- Rousseeuw, P.J. and Driessen, K.V., A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, 1999, **41**, 212–223.
- Ruder, S. and Plank, B., Strong baselines for neural semi-supervised learning under domain shift. *arXiv preprint arXiv:1804.09530*, 2018.
- Sharpe, W.F., Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 1964, **19**, 425–442.
- Song, K., Tan, X., Qin, T., Lu, J. and Liu, T.Y., MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 2020, **33**, 16857–16867.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F. *et al.*, Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Van Aelst, S. and Rousseeuw, P., Minimum volume ellipsoid. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2009, **1**, 71–82.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I., Attention is all you need. *Advances in neural information processing systems*, 2017, **30**.
- Wei, C., Fan, H., Xie, S., Wu, C.Y., Yuille, A. and Feichtenhofer, C., Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14668–14678, 2022.
- Welch, I. and Goyal, A., A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 2008, **21**, 1455–1508.
- Wen, Q., Zhang, Z., Li, Y. and Sun, L., Fast RobustSTL: Efficient and robust seasonal-trend decomposition for time series with complex patterns. In *Proceedings of the Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2203–2213, 2020.
- Woo, G., Liu, C., Sahoo, D., Kumar, A. and Hoi, S., Etsformer: Exponential smoothing transformers for time-series forecasting. *arXiv preprint arXiv:2202.01381*, 2022.
- Wu, H., Xu, J., Wang, J. and Long, M., Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 2021a, **34**, 22419–22430.
- Wu, H., Xu, J., Wang, J. and Long, M., Autoformer: Decomposition transformers with auto-correlation for

- long-term series forecasting. *Advances in Neural Information Processing Systems*, 2021b, **34**, 22419–22430.
- Xiao, T., Reed, C.J., Wang, X., Keutzer, K. and Darrell, T., Region similarity representation learning. In *Proceedings of the Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10539–10548, 2021.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R. and Le, Q.V., Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 2019, **32**.
- Yue, Z., Wang, Y., Duan, J., Yang, T., Huang, C., Tong, Y. and Xu, B., Ts2vec: Towards universal representation of time series. In *Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36, pp. 8980–8987, 2022.
- Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J. and Yoo, Y., Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6023–6032, 2019.
- Yun, S., Lee, H., Kim, J. and Shin, J., Patch-level representation learning for self-supervised vision transformers. In *Proceedings of the Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8354–8363, 2022.
- Zhai, X., Oliver, A., Kolesnikov, A. and Beyer, L., S4l: Self-supervised semi-supervised learning. In *Proceedings of the Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1476–1485, 2019.
- Zhang, H., Cisse, M., Dauphin, Y.N. and Lopez-Paz, D., mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Zhang, Z., Zohren, S. and Roberts, S., Deep learning for portfolio optimization. *arXiv preprint arXiv:2005.13665*, 2020.
- Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L. and Jin, R., Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *Proceedings of the International Conference on Machine Learning*, pp. 27268–27286, 2022.

Appendix

Appendix A: Pairs trading results for query stock and TOP@1 similar Stock

In this section, we present the pairs trading results obtained by applying the trading strategy to the query stock and its top@1 similar stock identified by SimStock and the benchmark models (TS2VEC, Corr1, Corr2, and Coint).

Table A.1 shows the terminal wealth and maximum drawdown (MDD) achieved for each query stock and its top@1 similar stock. The best-performing model for each query stock is highlighted in bold, while the second-best model is underlined.

SimStock achieves the highest terminal wealth for 3 out of the 12 query stocks (AAPL, CMG, and BA). Coint outperforms for 4 query stocks (V, XOM, AMZN, and MA), while Corr1 and Corr2 each outperform for 2 query stocks (MSFT and CVS for Corr1; PFE and MSFT for Corr2). TS2VEC achieves the highest wealth for AMZN. For the query stocks V and MA, SimStock, TS2VEC, Corr1, and Corr2 yield the same results.

In terms of maximum drawdown, SimStock achieves the lowest MDD for 4 out of the 12 query stocks (CMG, WFC, BA, and MSFT). TS2VEC exhibits the lowest MDD for 3 query stocks (AAPL, AMZN, and MA), while Corr1 outperforms for 2 query stocks (PFE and CVS). Corr2 achieves the lowest MDD for V and MA, and Coint outperforms for XOM.

TS2VEC fails to generate buy/sell signals for the query stock PFE and WFC, resulting in NaN values for wealth and MDD. Similarly, Corr1 fails to generate signals for one of the occurrences of the MSFT query stock. These instances show limitations in the consistency of the benchmark models in identifying tradable opportunities.

SimStock achieves competitive performance in terms of terminal wealth and MDD for a significant number of query stocks, indicating its effectiveness in capturing meaningful similarities that translate into profitable trading strategies.

Query Stock	Wealth					Maximum Drawdown (%)				
	SimStock	TS2VEC	Corr1	Corr2	Coint	SimStock	TS2VEC	Corr1	Corr2	Coint
AAPL	1508.74	1303.97	-798.23	767.28	NaN	-2.25	-0.46	-3.16	-11.6	NaN
CMG	373.21	-542.15	-1203.44	-1505.86	-59.09	<u>-5.68</u>	-4.92	-7.5	-24.78	-8.65
MSFT	804.26	1277.58	836.2	510.55	622.76	-3.82	<u>-5.42</u>	-5.55	-5.92	-11.76
WFC	<u>756.55</u>	NaN	1236.23	756.55	-2724.44	-0.05	<u>-5.74</u>	-6.82	-0.05	-31.55
V	<u>239.17</u>	<u>239.17</u>	<u>239.17</u>	<u>239.17</u>	701.17	0.00	0.00	0.00	0.00	<u>-2.78</u>
XOM	<u>104.92</u>	<u>104.92</u>	-192.2	-1120.8	1099.85	-11.41	-11.41	<u>-6.64</u>	-14.22	-1.40
PFE	-310.09	NaN	293.87	1501.43	440.50	-5.69	NaN	-1.88	-8.07	-7.04
AMZN	373.21	<u>1225.26</u>	-1743.54	661.29	2624.92	<u>-5.68</u>	-0.01	-22.64	-10.88	-27.18
BA	1321.94	117.8	-921.89	892.8	93.55	-3.66	-6.71	-11.96	-15.8	-22.80
MSFT	532.16	-2012.47	1619.35	1046.26	1789.72	-1.00	<u>-27.68</u>	NaN	-6.48	-26.11
MA	196.77	196.77	<u>196.77</u>	<u>196.77</u>	404.83	0.00	0.00	0.00	0.00	<u>-6.12</u>
CVS	726.03	-1464.39	1494.54	439.9	-1647.26	-2.58	<u>-9.36</u>	-10.24	-9.96	<u>-39.55</u>

Table A.1. Pairs trading results for the query stock and its top@1 similar stock identified by SimStock, TS2VEC, Corr1, Corr2 and Coint. The table presents the terminal wealth and maximum drawdown (MDD) achieved for each query stock and method combination. The best-performing method for each query stock is highlighted in bold, while the second-best method is underlined. NaN values indicate instances where the method failed to generate buy/sell signals.

Appendix B: Application to index tracking of thematic ETFs

Qualitative evaluation

In this subsection, to provide a comprehensive qualitative evaluation, we examine the overlapping and non-overlapping tickers for each model and ETF combination, along with their respective industries, as illustrated in Figure B.1. This analysis enables us to assess the extent to which the models capture the underlying themes of the ETFs. The experimental setup involves selecting 10 stocks to track each ETF on the US exchange, as shown in Figure 5.

For the ARKK ETF, which focuses on disruptive innovation, **SimStock** identifies several relevant stocks across various sectors. The overlapping tickers include ROKU (Roku Inc., a company that provides streaming media devices and services), SHOP (Shopify Inc., an e-commerce platform provider), and SQ (Square Inc., a financial technology company). The non-overlapping tickers selected by **SimStock** also align with the disruptive innovation theme, including companies from 3D printing (MTLS: Materialise NV; DDD: 3D Systems Corp.), biotechnology (EDIT: Editas Medicine Inc.), and e-commerce (MELI: MercadoLibre Inc.). In contrast, the baseline models (TS2VEC, and Corr1) either have no overlapping tickers or include stocks from less relevant industries, such as insurance (WTM: White Mountains Insurance Group Ltd.), sports (MANU: Manchester United plc), and real estate (FLJ: FLJ Group Ltd.).

For the SKYY ETF, which tracks the cloud computing industry, **SimStock** identifies NET (Cloudflare Inc., a content delivery network and DDoS mitigation company), cloud storage (DBX: Dropbox Inc.), financial software (INTU: Intuit Inc.), Elasticsearch (ESTC: Elastic N.V.) and PAYC (Paycom Software Inc., a provider of human capital management software) as overlapping tickers. The non-overlapping stocks selected by **SimStock** also show a strong connection to the cloud computing theme such as software compnay (TYL: Tyler Technologies Inc.; GWRE: Guidewire Software, Inc.; GLOB: Globant S.A.). The baseline models, however, include stocks from a wider range of industries that are less directly related to cloud computing, such as automotive (F: Ford Motor Co.), retail (W: Wayfair Inc.), and real estate (LAMR: Lamar Advertising Co.), biotechnology (LIFE: aTyr Pharma Inc.; VBIV: VBI Vaccines Inc.), and agriculture (VFF: Village Farms International Inc.).

For the BOTZ ETF, which focuses on robotics and artificial intelligence, **SimStock** identifies NVDA (NVIDIA Corp.) as an overlapping ticker. While the non-overlapping stocks selected by **SimStock** include some companies from other sectors, such as asset management (AMG: Affiliated Managers Group Inc.; GAM: General American Investors Co. Inc.) and building materials (MLM: Martin Marietta Materials Inc.), the model also identifies several companies from the semiconductor industry (QRVO: Qorvo Inc.; ASML: ASML Holding N.V.; LSCC: Lattice Semiconductor Corp.), which is a key enabler of robotics and AI. The baseline models, TS2VEC and Corr1, include stocks from a diverse range of industries that are less directly related to the robotics and AI theme, such as animal healthcare (IDXX: IDEXX Laboratories Inc.), insurance (AEG: Aegon N.V.), and electric vehicles (HYLN: Hyliion Holdings Corp.).

For the LIT ETF, which tracks the lithium and battery tech industry, **SimStock** identifies three overlapping tickers: ALB (Albemarle Corp.), ALTM (Altus Midstream Co.) and LAAC (Lithium Americas (Argentina) Corp), both of which are involved in lithium mining. However, the non-overlapping stocks selected by **SimStock** come from various other industries, such as asset management (AB: AllianceBernstein Holding L.P.; APAM: Artisan Partners Asset Management Inc.; BLK: BlackRock Inc.; KKR: KKR & Co. Inc.) and software (BB: BlackBerry Ltd.), which are less directly related to the lithium and battery tech theme. The baseline models also struggle to identify relevant stocks for this ETF, with no overlapping tickers and selections from industries such as insurance (WTM: White Mountains Insurance Group Ltd.; ERIE: Erie Indemnity Co.; ASR: Grupo Aeroportuario del Sureste, S. A. B. de C. V.), automotive (F: Ford Motor Co.; SUZ: Suzuki Motor Corp.), and engineering & construction (MYRG: MYR Group Inc.; FRHC: Freedom Holding Corp.).

This qualitative analysis represent that SimStock’s top@10 stocks for each ETF generally align better with the respective themes compared to the baseline models, particularly for the ARKK, SKYY, and BOTZ ETFs. However, the baseline models, TS2VEC and Corr1, include stocks from a diverse range of industries that are less directly related to the specific themes of the ETFs. Corr2 shows some improvement over the other baselines, with a few overlapping tickers for ARKK and BOTZ, but it still lags behind SimStock in terms of capturing the underlying themes of the ETFs.



Figure B.1. Top@10 stocks selected by each model to track the four thematic ETFs (ARKK, SKYY, BOTZ, and LIT) on the US exchange. The overlapping tickers between the model selections and the actual ETF holdings are shown in red, while the non-overlapping tickers are shown in white. SimStock demonstrates a higher number of overlapping tickers with the ETF holdings compared to the baseline models, particularly for ARKK, SKYY, and BOTZ, indicating its ability to identify stocks that align with the underlying themes of these ETFs.

Appendix C: Tracking Error Volatility (TEV) Results

In this section, we present the Tracking Error Volatility (TEV) results to complement the Tracking Error (TE) analysis discussed in the main text. TEV measures the volatility of the difference in returns between the tracking portfolio and the target ETF, providing information into the stability and consistency of the tracking performance.

We computed the TEV for each combination of the four thematic ETFs (ARK Innovation ETF (ARKK), First Trust Cloud Computing ETF (SKYY), Global X Robotics & AI ETF (BOTZ), and Global X Lithium & Battery Tech ETF (LIT)) and the four exchanges (US, SSE, SZSE, and TSE). The TEV values were calculated for the top k similar stocks ($k = 10, 15, 20, 25, 30, 35$) identified by `SimStock` and the baseline methods (Corr1, Corr2, TS2VEC).

Tables C.1 present the TEV results for ARKK, SKYY, BOTZ, and LIT, respectively. In each table, the best performing method for each k and exchange is highlighted in bold, while the second best is in underline. The TEV results are consistent with the findings from the TE analysis in the main text. `SimStock` achieves the lowest (bold) or second lowest (underline) TEV values for the majority of k values and exchanges, particularly for ARKK, SKYY, and BOTZ. This indicates that the similar stocks identified by `SimStock` not only closely track the performance of these ETFs but also exhibit lower volatility in the tracking error, resulting in more stable tracking portfolios.

The performance of `SimStock` in terms of both TE and TEV demonstrates its effectiveness in identifying stocks that can accurately and consistently track the performance of thematic ETFs, especially for ARKK, SKYY, and BOTZ. The lower TEV values achieved by `SimStock` suggest that its tracking portfolios are less prone to large deviations from the target ETF returns, providing investors with more reliable and stable tracking options.

ARK Innovation ETF (ARKK)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.1524	0.0893	0.0720	0.0630	0.0540	0.0474
	Corr2	0.1669	0.1141	0.1018	0.1061	0.0894	0.1042
	TS2VEC	0.1884	0.2225	0.1852	0.1701	0.1460	0.1292
	SimStock	0.1934	0.0764	0.0975	0.1785	0.1822	0.2414
SSE	Corr1	0.1747	0.1864	0.1820	0.1824	0.1833	0.1839
	Corr2	0.1513	0.1414	0.1417	0.1435	0.1394	0.1406
	TS2VEC	0.1633	0.1618	0.1527	0.1629	0.1637	0.1655
	SimStock	0.1414	0.1419	0.1352	0.1366	0.1358	0.1429
SZSE	Corr1	0.1747	0.1824	0.1709	0.1845	0.1800	0.1815
	Corr2	0.1399	0.1376	0.1355	0.1382	0.1394	0.1394
	TS2VEC	0.1495	0.1456	0.1489	0.1463	0.1410	0.1415
	SimStock	0.1436	0.1411	0.1410	0.1449	0.1397	0.1350
TSE	Corr1	0.1969	0.1759	0.1430	0.1388	0.1387	0.1358
	Corr2	0.1165	0.1129	0.1213	0.1172	0.1149	0.1141
	TS2VEC	0.1309	0.1335	0.1262	0.1279	0.1312	0.1251
	SimStock	0.1117	0.1116	0.1118	0.1125	0.1076	0.1074

First Trust Cloud Computing ETF (SKYY)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.1150	0.1406	0.1379	0.1245	0.1058	0.0993
	Corr2	0.2410	0.2298	0.2011	0.1916	0.1586	0.1455
	TS2VEC	0.1849	0.1655	0.1407	0.1397	0.1387	0.1330
	SimStock	0.0616	0.0364	0.0561	0.0653	0.0541	0.0521
SSE	Corr1	0.1699	0.1929	0.1993	0.2014	0.2035	0.2134
	Corr2	0.1490	0.1386	0.1303	0.1118	0.1133	0.1188
	TS2VEC	0.1897	0.1975	0.2151	0.2074	0.2013	0.1871
	SimStock	0.1106	0.0921	0.1128	0.1072	0.1074	0.1167
SZSE	Corr1	0.2248	0.2183	0.2123	0.2156	0.2106	0.2042
	Corr2	0.1103	0.1143	0.1164	0.1122	0.1212	0.1239
	TS2VEC	0.1642	0.1796	0.1684	0.1689	0.1702	0.1594
	SimStock	0.1103	0.1130	0.1157	0.1105	0.1047	0.1020
TSE	Corr1	0.1308	0.1363	0.1239	0.1429	0.1450	0.1472
	Corr2	0.0723	0.0846	0.0762	0.0871	0.0818	0.0810
	TS2VEC	0.1066	0.1142	0.1153	0.1136	0.1108	0.1215
	SimStock	0.0842	0.0703	0.0704	0.0640	0.0570	0.0625

Global X Robotics & AI ETF (BOTZ)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.1599	0.1030	0.0744	0.0758	0.0729	0.0774
	Corr2	0.0807	0.0971	0.0872	0.0892	0.0951	0.0865
	TS2VEC	0.1023	0.1464	0.1458	0.1218	0.1221	0.1338
	SimStock	0.0359	0.0360	0.0465	0.0535	0.0551	0.0548
SSE	Corr1	0.1326	0.1354	0.1379	0.1337	0.1315	0.1348
	Corr2	0.0810	0.0846	0.0814	0.0819	0.0846	0.0887
	TS2VEC	0.1274	0.1204	0.1196	0.1315	0.1251	0.1225
	SimStock	0.0664	0.0745	0.0760	0.0828	0.0782	0.0755
SZSE	Corr1	0.1236	0.1143	0.1193	0.1250	0.1333	0.1369
	Corr2	0.0910	0.0884	0.0868	0.0836	0.0837	0.0808
	TS2VEC	0.0914	0.0812	0.0855	0.0815	0.0804	0.0890
	SimStock	0.0808	0.0875	0.1083	0.1030	0.1044	0.0991
TSE	Corr1	0.0688	0.0680	0.0729	0.0757	0.0724	0.0643
	Corr2	0.0773	0.0890	0.0837	0.0830	0.0864	0.0818
	TS2VEC	0.0970	0.0926	0.0834	0.0861	0.0804	0.0814
	SimStock	0.0604	0.0730	0.0711	0.0792	0.0794	0.0764

Global X Lithium & Battery Tech ETF (LIT)							
Exchange	Methods	TOP@10	TOP@15	TOP@20	TOP@25	TOP@30	TOP@35
US	Corr1	0.1548	0.1744	0.1605	0.1708	0.1805	0.1706
	Corr2	0.0935	0.1440	0.1450	0.1433	0.1677	0.1246
	TS2VEC	0.0819	0.1091	0.0858	0.0869	0.1051	0.0996
	SimStock	0.0389	0.0513	0.0501	0.0492	0.0560	0.0536
SSE	Corr1	0.0840	0.0752	0.0729	0.0701	0.0724	0.0714
	Corr2	0.0843	0.0837	0.0804	0.0843	0.0883	0.0901
	TS2VEC	0.0702	0.0781	0.0764	0.0788	0.0817	0.0803
	SimStock	0.0945	0.1193	0.1031	0.0942	0.1060	0.1114
SZSE	Corr1	0.0565	0.0567	0.0617	0.0609	0.0622	0.0679
	Corr2	0.2148	0.1898	0.1643	0.1509	0.1512	0.1399
	TS2VEC	0.0629	0.0661	0.0704	0.0723	0.0773	0.0821
	SimStock	0.2200	0.2361	0.2102	0.2057	0.1963	0.1817
TSE	Corr1	0.1522	0.1718	0.1704	0.1839	0.1713	0.1556
	Corr2	0.1768	0.1807	0.1603	0.1575	0.1851	0.1771
	TS2VEC	0.1488	0.1372	0.1444	0.1447	0.1318	0.1426
	SimStock	0.1580	0.1750	0.1715	0.1808	0.1837	0.1772

Table C.1. Tracking errors volatility (TEV) of SimStock and baseline methods for tracking the performance of four thematic ETFs (ARKK, SKYY, BOTZ, and LIT) using top k similar stocks from the US, SSE, SZSE, and TSE exchanges. The best performing method for each k and exchange is highlighted in black, while the second best is in underscore.

Appendix D: Application to Portfolio optimization

Performance Statistics of Major Stock Indices

In this section, we presents the performance statistics of four major stock indices: S&P 500, JPX Prime 150, SSE 100, and SZSE 100 in table D.1. The evaluation period spans from January 2023 to Feb 2024. The S&P 500 and JPX Prime 150 indices exhibited positive returns during this period, while the SSE 100 and SZSE 100 indices experienced negative returns.

The S&P 500 index, which represents the U.S. stock market, generated an arithmetic return of 4.87% and a geometric return of 2.71%. The cumulative return over the evaluation period was 13.02%. The annualized standard deviation of 30.84% indicates the high volatility of the index. The maximum drawdown of -31.32% highlights the substantial decline the index experienced during the period. The monthly 95% Value at Risk (VaR) of -11.9% indicates the potential loss that could be incurred with a 95% confidence level. The Sharpe ratio of 0.06 suggests a low risk-adjusted return for the S&P 500 index.

The JPX Prime 150 index, representing the Japanese stock market, outperformed the other indices with an arithmetic return of 13.69% and a geometric return of 11.61%. The cumulative return over the evaluation period was 49.59%. The annualized standard deviation of 27.18% was lower than that of the S&P 500 index, indicating relatively lower volatility. The annualized skewness and kurtosis values of -0.04 and 2.89, respectively, were similar to those of the S&P 500 index. The maximum drawdown of -23.73% was less severe compared to the other indices. The monthly 95% VaR of -9.94% was also lower than that of the S&P 500 index. The Sharpe ratio of 0.58 indicates a higher risk-adjusted return for the JPX Prime 150 index compared to the S&P 500 index.

The SSE 100 and SZSE 100 indices, representing the Chinese stock market, experienced negative returns during the evaluation period. The SSE 100 index had an arithmetic return of -5.36% and a geometric return of -6.99%, resulting in a cumulative return of -11.16%. The SZSE 100 index had an arithmetic return of -7.81% and a geometric return of -9.86%, resulting in a cumulative return of -21.02%. Both indices exhibited higher volatility compared to the S&P 500 and JPX Prime 150 indices, with annualized standard deviations of 39.13% and 37.39%, respectively. The annualized skewness and kurtosis values were close to zero and slightly higher than three, respectively, indicating a relatively symmetric return distribution with slightly heavier tails compared to a normal distribution. The maximum drawdowns of -45.08% and -46.1% for the SSE 100 and SZSE 100 indices, respectively, were more severe compared to the S&P 500 and JPX Prime 150 indices. The monthly 95% VaR values of -16.24% and -16.16% were also higher, indicating a higher potential loss. The Sharpe ratios of -0.35 and -0.48 for the SSE 100 and SZSE 100 indices, respectively, suggest poor risk-adjusted returns during the evaluation period.

Evaluation Method	Exchange			
	S&P500	JPX Prime 150	SSE 100	SZSE 100
Covariance Method				
Arithmetic Return (%)	4.87	13.69	-5.36	-7.81
Geometric Return (%)	2.71	11.61	-6.99	-9.86
Cumulative Return (%)	13.02	49.59	-11.16	-21.02
Annualized SD (%)	30.84	27.18	39.13	37.39
Annualized Skewness	-0.03	-0.04	0.04	0.01
Annualized Kurtosis	2.81	2.89	3.12	3.24
Maximum Drawdown (%)	-31.32	-23.73	-45.08	-46.1
Monthly 95% VaR (%)	-11.9	-9.94	-16.24	-16.16
Sharpe Ratio	0.06	0.58	-0.35	-0.48

Table D.1. Performance statistics of major stock indices from January 2023 to February 2024. The table presents various return and risk measures for the S&P 500, JPX Prime 150, SSE 100, and SZSE 100 indices. The S&P 500 and JPX Prime 150 indices exhibited positive returns, while the SSE 100 and SZSE 100 indices experienced negative returns during the evaluation period.

Performance Metrics for Portfolio Optimization with 10 and 50 Stocks

This appendix presents the performance metrics for MVO with 10 and 50 stocks in the S&P 500 and JPX Prime 150 universes. The four portfolio construction methods evaluated are Simstock Embedding (SS), Historical Covariance (HC), Shrinkage Method (SM), and Gerber Statistic (GS).

We investigate the performance of the four covariance estimation methods (SS, HC, SM, GS, TS) for constructing portfolios with 10 and 50 stocks in the S&P 500 and JPX Prime 150 universes at different target risk levels (24%, 27%, 30%, and 33%).

Tables C.3 and D.3 present the performance metrics for 10-stock portfolios in the S&P 500 and JPX Prime 150, respectively. In both universes, the SS method generally outperforms the other methods across all risk target levels, achieving higher returns while maintaining similar risk measures. As the risk target level increases, the returns and risk measures for all methods also increase.

Tables C.5 and D.5 show the performance metrics for 50-stock portfolios in the S&P 500 and JPX Prime 150, respectively. Similar to the 10-stock portfolios, the SS method demonstrates superior performance compared to the other methods across all risk target levels in both universes. The HC method has the second-best performance, followed by the SM, GS and TS methods. As the risk target level increases, the returns and risk measures for all methods increase.

The consistent outperformance of the Simstock Embedding (SS) method in terms of returns for both 10-stock and 50-stock portfolios in the S&P 500 and JPX Prime 150 universes highlights its potential as an effective stock embedding technique for portfolio optimization.

S&P500 — 10 Stocks	Target Volatility (24%)				Target Volatility (27%)				Target Volatility (30%)				Target Volatility (33%)							
	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS
Covariance Method	7.94	6.85	6.84	6.97	5.28	8.38	7.35	7.27	7.27	6.02	8.56	7.71	7.8	7.71	7.63	8.84	8.12	8.27	8.13	7.63
Arithmetic Return (%)	6.78	6.05	6.08	6.11	3.63	7.1	6.41	6.34	6.25	4.84	7.18	6.61	6.68	6.51	6.5	7.31	6.85	6.98	6.77	6.50
Geometric Return (%)	24.63	22.11	22.11	22.32	11.29	26.26	23.73	23.49	23.19	15.24	26.79	24.73	25.05	24.39	20.78	27.52	25.71	26.3	25.53	20.78
Cumulative Return (%)	27.22	24.68	24.57	25.06	27.33	28.18	26.06	26.11	26.49	28.32	29.11	27.31	27.47	27.79	29.55	29.94	28.44	28.67	28.92	29.55
Annualized SD (%)	-0.06	-0.06	-0.03	-0.03	-0.09	-0.08	-0.07	-0.06	-0.06	-0.11	-0.09	-0.08	-0.07	-0.08	-0.12	-0.10	-0.08	-0.08	-0.08	-0.12
Annualized Skewness	3.04	3.06	3.04	3.04	2.98	3.09	3.08	3.07	3.07	2.95	3.12	3.12	3.11	3.11	2.93	3.15	3.15	3.15	3.15	2.93
Annualized Kurtosis	-26.07	-24.22	-24.04	-24.4	-26.91	-27.13	-25.28	-25.23	-25.73	-26.75	-27.93	-26.52	-26.41	-26.93	-26.27	-28.59	-27.49	-27.46	-27.85	-26.27
Maximum Drawdown (%)	-10.24	-9.43	-9.39	-9.57	-10.9	-10.62	-9.92	-9.93	-10.06	-11.21	-11.00	-10.30	-10.36	-10.49	-11.54	-11.28	-10.7	-10.76	-10.85	-11.54
Monthly 95% VaR (%)	0.24	0.22	0.23	0.23	0.07	0.25	0.23	0.23	0.22	0.16	0.25	0.23	0.23	0.22	0.24	0.24	0.23	0.24	0.23	0.24
Sharpe Ratio	7.15	6.81	6.88	6.94	6.54	7.31	7.00	7.09	7.15	6.72	7.22	7.15	7.21	7.25	6.90	7.23	7.21	7.24	7.26	6.90
Annualized Turnover																				

Table D.2. This table presents the performance metrics for four portfolio construction methods in the JPX Prime 150: Simstock embedding(SS), historical covariance (HC), shrinkage method (SM), and Gerber statistic (GS). The portfolios were optimized for four different risk target levels: 24%, 27%, 30%, and 33%. The performance was evaluated over the full testing period from January 2022 to February 2024. The 3-month U.S. Treasury Bill rate was used as the risk-free rate for performance calculations. Transaction costs were modeled as 10 basis points of the traded volume for each rebalancing event.

S&P500 — 10 Stocks	Target Volatility (24%)				Target Volatility (27%)				Target Volatility (30%)				Target Volatility (33%)							
	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS
Covariance Method	7.94	6.85	6.84	6.97	5.28	8.38	7.35	7.27	7.27	6.02	8.56	7.71	7.8	7.71	7.63	8.84	8.12	8.27	8.13	7.63
Arithmetic Return (%)	6.78	6.05	6.08	6.11	3.63	7.1	6.41	6.34	6.25	4.84	7.18	6.61	6.68	6.51	6.5	7.31	6.85	6.98	6.77	6.50
Geometric Return (%)	24.63	22.11	22.11	22.32	11.29	26.26	23.73	23.49	23.19	15.24	26.79	24.73	25.05	24.39	20.78	27.52	25.71	26.3	25.53	20.78
Cumulative Return (%)	27.22	24.68	24.57	25.06	27.33	28.18	26.06	26.11	26.49	28.32	29.11	27.31	27.47	27.79	29.55	29.94	28.44	28.67	28.92	29.55
Annualized SD (%)	-0.06	-0.06	-0.03	-0.03	-0.09	-0.08	-0.07	-0.06	-0.06	-0.11	-0.09	-0.08	-0.07	-0.08	-0.12	-0.10	-0.08	-0.08	-0.08	-0.12
Annualized Skewness	3.04	3.06	3.04	3.04	2.98	3.09	3.08	3.07	3.07	2.95	3.12	3.12	3.11	3.11	2.93	3.15	3.15	3.15	3.15	2.93
Annualized Kurtosis	-26.07	-24.22	-24.04	-24.4	-26.91	-27.13	-25.28	-25.23	-25.73	-26.75	-27.93	-26.52	-26.41	-26.93	-26.27	-28.59	-27.49	-27.46	-27.85	-26.27
Maximum Drawdown (%)	-10.24	-9.43	-9.39	-9.57	-10.9	-10.62	-9.92	-9.93	-10.06	-11.21	-11.00	-10.30	-10.36	-10.49	-11.54	-11.28	-10.7	-10.76	-10.85	-11.54
Monthly 95% VaR (%)	0.24	0.22	0.23	0.23	0.07	0.25	0.23	0.23	0.22	0.16	0.25	0.23	0.23	0.22	0.24	0.24	0.23	0.24	0.23	0.24
Sharpe Ratio	7.15	6.81	6.88	6.94	6.54	7.31	7.00	7.09	7.15	6.72	7.22	7.15	7.21	7.25	6.90	7.23	7.21	7.24	7.26	6.90
Annualized Turnover																				

Table D.3. This table presents the performance metrics for four portfolio construction methods in the JPX Prime 150: Simstock embedding(SS), historical covariance (HC), shrinkage method (SM), and Gerber statistic (GS). The portfolios were optimized for four different risk target levels: 24%, 27%, 30%, and 33%. The performance was evaluated over the full testing period from January 2022 to February 2024. The 3-month U.S. Treasury Bill rate was used as the risk-free rate for performance calculations. Transaction costs were modeled as 10 basis points of the traded volume for each rebalancing event.

S&P500 — 50 Stocks	Target Volatility (24%)					Target Volatility (27%)					Target Volatility (30%)					Target Volatility (33%)				
	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS
Covariance Method	10.41	10.27	9.82	10.54	10.65	11.86	11.15	11.38	11.59	11.63	12.82	12.1	12.7	12.67	12.66	13.87	13.03	13.78	13.57	13.17
Arithmetic Return (%)	9.59	9.5	9.19	9.84	8.94	10.92	10.24	10.57	10.71	10.39	11.74	11.01	11.66	11.56	11.26	12.57	11.73	12.47	12.27	11.32
Geometric Return (%)	34.15	33.49	32.06	34.49	29.29	39.14	36.27	37.5	38.06	34.54	42.3	39.41	41.93	41.71	37.75	45.6	42.45	45.44	44.76	37.95
Cumulative Return (%)	29	27.93	27.2	27.53	30.31	30.6	29.67	29.42	29.58	31.09	32.34	31.38	31.54	31.64	32.28	33.92	33.07	33.47	33.52	33.13
Annualized SD (%)	-0.17	-0.22	-0.19	-0.18	-0.17	-0.2	-0.25	-0.24	-0.23	-0.21	-0.23	-0.27	-0.27	-0.26	-0.19	-0.26	-0.29	-0.31	-0.29	-0.19
Annualized Skewness	3.12	3.28	3.32	3.19	2.83	3.18	3.35	3.36	3.27	2.84	3.24	3.37	3.4	3.31	2.87	3.24	3.36	3.38	3.3	2.97
Annualized Kurtosis	-25.4	-24.36	-23.35	-23.44	-26.04	-25.79	-25.44	-24.66	-24.97	-26.18	-26.33	-26.6	-25.94	-26.53	-26.15	-27.29	-27.68	-27.34	-27.85	-26.57
Maximum Drawdown (%)	-10.96	-10.36	-10.15	-10.18	-11.38	-11.28	-10.79	-10.78	-10.82	-11.83	-11.64	-11.2	-11.3	-11.29	-12.22	-12.06	-11.68	-11.8	-11.79	-12.40
Monthly 95% VaR (%)	0.39	0.41	0.39	0.43	0.36	0.44	0.42	0.44	0.45	0.40	0.46	0.43	0.46	0.45	0.48	0.44	0.47	0.47	0.46	
Sharpe Ratio	8.94	8.57	8.69	8.5	8.10	8.9	8.63	8.74	8.57	8.21	8.84	8.66	8.67	8.6	8.04	8.7	8.63	8.55	8.54	
Annualized Turnover																				

Table D.4. This table presents the performance metrics for four portfolio construction methods in the JPX Prime 150: Simstock embedding(SS), historical covariance (HC), shrinkage method (SM), and Gerber statistic (GS). The portfolios were optimized for four different risk target levels: 24%, 27%, 30%, and 33%. The performance was evaluated over the full testing period from January 2022 to February 2024. The 3-month U.S. Treasury Bill rate was used as the risk-free rate for performance calculations. Transaction costs were modeled as 10 basis points of the traded volume for each rebalancing event.

JPX Prime 150 — 50 Stocks	Target Volatility (24%)					Target Volatility (27%)					Target Volatility (30%)					Target Volatility (33%)				
	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS	SS	HC	SM	GS	TS
Covariance Method	20.3	19.7	18.6	17.55	15.69	21.3	20.03	19.79	18.99	15.32	22.74	20.34	21.2	20.22	15.46	23.49	20.93	21.93	21.25	15.84
Arithmetic Return (%)	18.9	18.52	17.31	16.51	14.00	19.89	18.79	18.38	17.8	14.16	21.27	19.04	19.68	18.96	13.82	21.98	19.52	20.34	19.86	14.74
Geometric Return (%)	70.97	69.2	63.79	59.9	48.15	75.61	70.44	68.66	65.51	48.77	82.1	71.7	74.76	70.89	47.44	85.94	74.21	78.27	75.42	51.04
Cumulative Return (%)	28.13	27.37	26.14	26.54	26.85	29.89	28.77	28.08	28.34	28.32	31.52	30.25	29.98	30.02	29.28	33.01	31.65	31.42	31.43	30.25
Annualized SD (%)	-0.03	0.00	-0.01	-0.01	-0.10	0.02	-0.03	0.02	0.01	-0.09	0.03	-0.03	0.02	-0.01	-0.08	0.03	-0.06	-0.02	-0.07	-0.03
Annualized Skewness	2.89	2.86	2.83	2.79	2.62	2.88	2.88	2.90	2.87	2.66	2.89	2.88	2.93	2.88	2.63	2.9	2.91	2.91	2.91	2.62
Annualized Kurtosis	-21.47	-20.57	-20.47	-20.64	-23.59	-22.41	-21.99	-21.58	-21.61	-23.34	-22.99	-23.12	-22.54	-22.45	-24.94	-23.79	-24.39	-23.67	-23.67	-25.09
Maximum Drawdown (%)	-9.73	-9.68	-8.9	-9.27	-10.17	-10.25	-10.23	-9.55	-9.87	-10.59	-10.74	-10.8	-10.22	-10.51	-10.99	-11.22	-11.35	-10.83	-11.18	-11.51
Monthly 95% VaR (%)	0.93	0.93	0.91	0.85	0.68	0.93	0.9	0.91	0.86	0.64	0.96	0.87	0.92	0.88	0.62	0.95	0.85	0.91	0.88	0.62
Sharpe Ratio	8.99	8.77	8.98	8.96	8.67	9.00	8.91	9.08	8.99	8.57	8.91	9.01	9.00	8.93	8.76	8.86	8.98	8.97	8.90	8.80
Annualized Turnover																				

Table D.5. This table presents the performance metrics for four portfolio construction methods in the JPX Prime 150: Simstock embedding(SS), historical covariance (HC), shrinkage method (SM), and Gerber statistic (GS). The portfolios were optimized for four different risk target levels: 24%, 27%, 30%, and 33%. The performance was evaluated over the full testing period from January 2022 to February 2024. The 3-month U.S. Treasury Bill rate was used as the risk-free rate for performance calculations. Transaction costs were modeled as 10 basis points of the traded volume for each rebalancing event.