RISK-AWARE STOCK RECOMMENDATION VIA FUNCTIONAL CAUSAL MODELING

Anonymous authors
Paper under double-blind review

ABSTRACT

Stock recommendation aims to rank stocks by their expected returns based on historical market data. Most existing methods model inter-stock correlations, which may fail to capture the directed information flow induced by asymmetric relations such as supplier-consumer dependencies. They also often prioritize profitability without explicitly accounting for stock-specific risk. To address these limitations, we propose DISCO, a unified risk-aware stock recommendation framework that combines temporal causal discovery with functional causal modeling. DISCO learns directed relations among stocks and models dynamic risk as a stochastic noise variable, enabling it to jointly estimate future returns and stock-level uncertainty. An optional multi-situation inference strategy further accounts for uncertainty in both the learned causal graph and stock risk. Experiments on three real-world stock markets demonstrate that DISCO consistently outperforms strong baselines in profitability and risk-adjusted performance. Code is available at https://anonymous.4open.science/r/DISCO-D088.

1 INTRODUCTION

Stock recommendation aims to rank stocks by their expected returns and plays an important role in investment decision support and automated trading Sanz-Cruzado et al. (2022). Unlike conventional recommendation tasks that personalize results using user interaction histories, stock recommendation typically operates without access to individual transaction data because of privacy and regulatory constraints Hsu et al. (2021). It therefore focuses on producing a market-level ranking of stocks from historical price observations. Existing methods commonly adopt an encoder–decoder architecture to model market dynamics, predict future returns, and rank stocks accordingly Hsu et al. (2021); Chen et al. (2025). However, the high volatility and complex dependencies of financial markets make reliable return prediction particularly challenging Sawhney et al. (2021).

Since the return of one stock may depend on other stocks in the same market, existing methods commonly model inter-stock correlations using attention mechanisms Wang et al. (2022); Ma & Tan (2022) or graph learning Sawhney et al. (2021); Feng et al. (2022). However, many economically meaningful relationships, such as supplier–consumer dependencies, are directional rather than symmetric Feng et al. (2019b). Correlation-based models may therefore capture predictive associations without identifying the underlying direction of information flow. For example, consider three stocks satisfying $X^2 = 0.5X^1 + z^2, X^3 = X^2 + z^3$, where $z^2, z^3 \sim \mathrm{Uniform}(0, 1)$ Goudet et al. (2018). Although the true structure is $X^1 \to X^2 \to X^3$, correlation-based regression may use X^3 to predict X^2 and thus introduce a spurious reverse dependency. Causal discovery instead seeks to recover the directed dependencies that generate these observations, providing a more faithful information flow among stocks. Nevertheless, causal structure remains underexplored in stock recommendation.

Risk awareness presents another important challenge. Most stock recommendation methods primarily optimize expected returns Wang et al. (2022); Zheng et al. (2023), potentially favoring stocks with high predicted profitability but substantial uncertainty. Although risk has been extensively studied in financial modeling Cheng et al. (2023a); Luo et al. (2023), existing methods often quantify it using static measures such as historical return variance Luo et al. (2023); Wang et al. (2019). In practice, stock risk evolves with market conditions and may propagate through inter-stock dependencies Cheng & Li (2022); MYRON & Scholes (2000). Accurately ranking stocks therefore requires jointly

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

modeling expected returns and dynamic, stock-specific risk, which remains underexplored in existing recommendation methods.

To address these limitations, we propose **DISCO**, a unified stock recommendation framework integrating **D**ynamic r**IS**k modeling and **C**ausal disc**O**very. DISCO first learns a posterior distribution over temporal causal graphs to characterize directed dependencies among stocks. It then incorporates the discovered causal structure into a functional causal model (FCM) Gong et al. (2022), which decomposes each stock's future return into market-driven information and a stochastic noise variable representing stock-specific risk. This formulation enables DISCO to jointly model directed market dependencies, future returns, and risk uncertainty. To further account for uncertainty in both the causal graph and stock risk, we introduce an optional multi-situation inference strategy that aggregates predictions over multiple graph and risk samples. Our main contributions are as follows:

- We introduce a causal stock recommendation framework that learns directed temporal dependencies among stocks, providing a faithful representation of market information flow.
- We develop an FCM-based return model that represents stock-specific risk as a stochastic component separated from market-driven effects, enabling joint return prediction and risk modeling.
- We propose a multi-situation inference strategy that accounts for uncertainty in both market dependencies and stock-specific risk. Experiments on three real-world stock markets demonstrate the effectiveness of DISCO in terms of profitability and risk-adjusted performance.

2 Related Work

2.1 Stock Relation Discovery

Existing studies model stock relations using either prior inter-company knowledge or correlations learned from market data. Some methods Feng et al. (2019b); Hsu et al. (2021) use explicit relations such as "investing " and "subsidiary". However, such relations may be inaccessible because of confidentiality and privacy concerns, and can become outdated as companies evolve. Consequently, recent studies increasingly learn stock relations directly from data.

Correlation-based Relation Modeling. Graph Neural Networks and Graph Convolutional Networks are widely used to model stock correlations Feng et al. (2019b); Hsu et al. (2021); Wang et al. (2022); Zheng et al. (2023); Cheng et al. (2023b); Ma et al. (2022); Lei et al. (2024). Following the success of Transformers in sequential modeling Vaswani et al. (2017), attention mechanisms have also been applied to stock relation learning Sawhney et al. (2021); Wang et al. (2022).

Causal Discovery. Causal discovery aims to identify causal relations from observational data Glymour et al. (2019). Methods for panel and time-series data generally fall into three categories: (1) constraint-based methods, including PC Spirtes & Glymour (1991), FCI Spirtes et al. (2000), and PCMCI Gerhardus & Runge (2020), which infer causal graphs through conditional independence tests; (2) score-based methods, which optimize statistical scores using techniques such as penalized neural ordinary differential equations and acyclicity constraints Pamfil et al. (2020); Zheng et al. (2018); and (3) functional causal models (FCMs) Shimizu et al. (2006); Hoyer et al. (2008), which represent causal relations as directed acyclic graphs and model causal mechanisms using nonlinear functions such as neural networks Gong et al. (2022); Geffner et al. (2022). Within the FCM family, DECI Geffner et al. (2022) is an end-to-end flow-based nonlinear additive-noise model for observational data, while Rhino Gong et al. (2022) extends this framework to time-series causal discovery. We adopt an FCM to learn inter-stock causal relations within an end-to-end recommendation framework.

2.2 Stock Risk Modeling

Reliable stock investment requires balancing returns and risks, particularly in leveraged trading Luo et al. (2023). Although risk modeling remains relatively underexplored in stock recommendation, it has long been studied in finance. The Markowitz mean-variance model Markowitz (1991) represents expected profit by mean return and risk by return variance. Subsequent studies have sought to balance these objectives Luo et al. (2023); Sun et al. (2023), with Luo et al. (2023) incorporating price volatility directly into the training objective. However, such static measures do not capture the evolving nature of risk in real markets Luo et al. (2023); Wang et al. (2019). To model risk more

adaptively, Cheng et al. (2023b) introduces a perturbation generator that increases model sensitivity to adversarial perturbations in high-risk scenarios. Nevertheless, dynamic risk modeling remains insufficiently explored in quantitative finance.

3 PRELIMINARY

3.1 TASK FORMULATION

Consider a market containing D stocks and a target trading day t. The input consists of price features from the preceding L trading days, denoted by $\boldsymbol{X}_{<t} = [\boldsymbol{X}_{t-1}, \boldsymbol{X}_{t-2}, \ldots, \boldsymbol{X}_{t-L}] = \{X^i_{t-l}\}^{D,L}_{i=1,l=1} \in \mathbb{R}^{D\times L\times d}$. Here, $X^i_{t-l} \in \mathbb{R}^d$ denotes the feature vector of stock i on day $t-l$. Following Feng et al. (2019b), the features include the normalized closing price p^i_{t-l} and its moving averages over 5, 10, 20, and 30 trading days, capturing trends at different temporal scales. Given $\boldsymbol{X}_{<t}$, the task is to predict the one-day return ratios $\boldsymbol{r}_t = \{r^i_t\}^D_{i=1}$ of all stocks on day t, where $\boldsymbol{r}_t = (\boldsymbol{p}_t - \boldsymbol{p}_{t-1})/\boldsymbol{p}_{t-1}$ and $\boldsymbol{p}_t = \{p^i_t\}^D_{i=1} \in \mathbb{R}^D$. The division is performed element-wise. Stocks are then ranked in descending order of their predicted returns, and the highest-ranked stocks are recommended. Equivalently, the task can be formulated as learning the conditional distribution $p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t})$. We adopt a generative framework that first discovers causal relations among stocks and then incorporates the learned relations and dynamic risk into an FCM to generate future returns.

3.2 TEMPORAL CAUSAL GRAPH & FUNCTIONAL CAUSAL MODEL (FCM)

Consider a multivariate time series $\{X^i_t\}^D_{i=1}$ with D variables and a maximum time lag L. A temporal causal graph is represented by a collection of lag-specific adjacency matrices $\boldsymbol{G} = [G_1, G_2, \ldots, G_L] \in \{0,1\}^{L\times D\times D}$, where $G^{ji}_l = 1$ indicates a directed causal relation from X^j_{t-l} to X^i_t, and $G^{ji}_l = 0$ otherwise Pamfil et al. (2020). Since all edges point from historical to current variables, the temporal graph is naturally acyclic. Figure 1(a) illustrates a causal mechanism with $D = 3$ variables and $L = 3$ lags, while Figure 1(b) represents the same mechanism using the adjacency matrices G_1, G_2, and G_3.

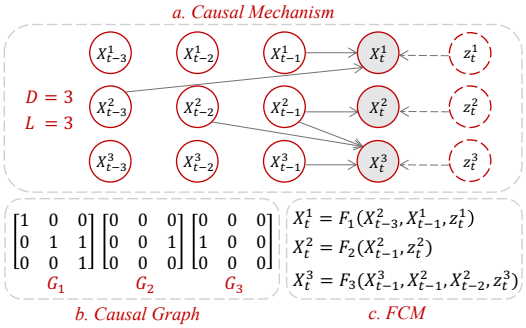

Figure 1: An FCM case for temporal data X^i_t with $D = 3$ variables and time lag $L = 3$. (a) A causal mechanism example. (b) The causal graph according to the causal mechanism. (c) The FCM based on the causal graph.

A functional causal model (FCM), also known as a structural equation model (SEM), specifies how each variable is generated from its causal parents and an exogenous noise term Gong et al. (2022). Given $\boldsymbol{G}$, the generating mechanism for variable X^i_t is written as $X^i_t = F_i\big(\mathbf{Pa}^i_{\boldsymbol{G}}(<t), z^i_t\big)$, where $\mathbf{Pa}^i_{\boldsymbol{G}}(<t) = \{X^j_{t-l} \mid G^{ji}_l = 1\}$ denotes the historical parents of X^i_t, and z^i_t captures variation not explained by these parents. The noise terms are assumed to be mutually independent across variables, $z^i_t \perp\!\!\!\perp z^j_t$ for $i \neq j$, serially independent over time, $z^i_t \perp\!\!\!\perp z^i_{t'}$ for $t \neq t'$, and independent of the historical variables. The function F_i characterizes how the parents and noise jointly generate X^i_t. Figure 1(c) shows the FCM equations induced by the temporal causal graph in Figure 1(b), illustrating the correspondence between graph edges and functional dependencies.

In general, different causal graphs and functional mechanisms may induce the same observational distribution, leading to causal structure non-identifiability Zhang et al. (2015). A common way to improve identifiability is to restrict the functional class. In this work, we consider additive-noise FCMs Hoyer et al. (2008); Shimizu et al. (2006), where $X^i_t = F_i\big(\mathbf{Pa}^i_{\boldsymbol{G}}(<t), z^i_t\big) = f_i\big(\mathbf{Pa}^i_{\boldsymbol{G}}(<t)\big) + z^i_t$. Here, f_i models the effects of the causal parents, while z^i_t is separated from the causal mechanism. Learning the graph $\boldsymbol{G}$, the functions f_i, and the noise distributions jointly specifies the generative distribution of the multivariate time series.

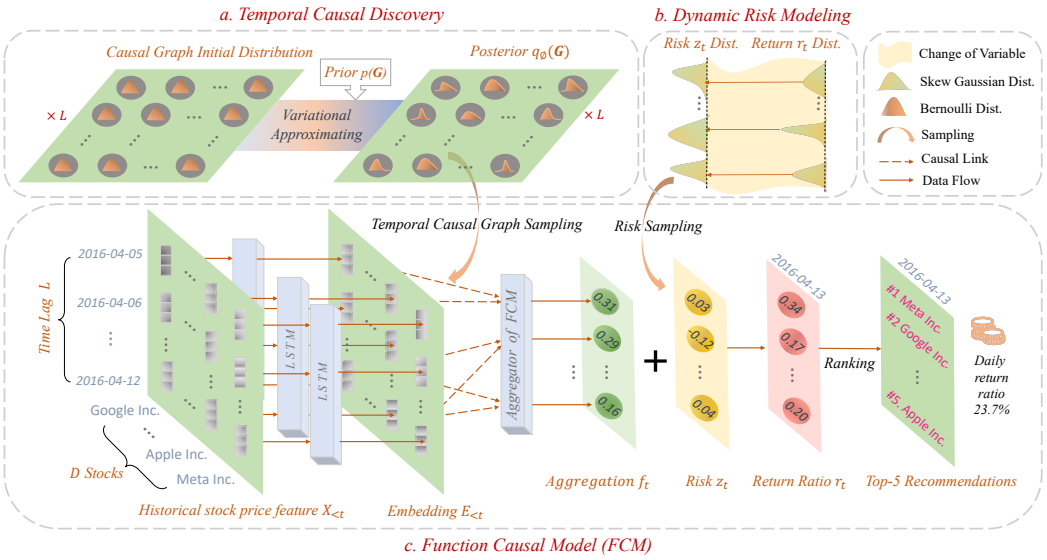

Figure 2: Overview of DISCO. (a) Temporal causal discovery, where variational inference combines the graph prior with a factorized Bernoulli approximation to learn the posterior distribution over causal graphs. (b) Dynamic risk modeling, where the change-of-variables formula relates the stochastic risk distribution to the conditional return distribution. (c) Risk-aware return generation, where the FCM aggregates historical information according to a sampled causal graph, incorporates sampled stock-specific risk, and ranks stocks by their predicted returns.

4 Method: DISCO

4.1 Model Overview

Given the historical market observations $\boldsymbol{X}_{<t}$, DISCO predicts the future return vector $\boldsymbol{r}_t$ and ranks stocks accordingly. Rather than assuming a fixed relation structure, DISCO treats the temporal causal graph $\boldsymbol{G}$ as a latent variable and marginalizes over its possible structures. The predictive distribution is therefore formulated as $p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t}) = \int_{\boldsymbol{G}} p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t}, \boldsymbol{G}) p(\boldsymbol{G} \mid \boldsymbol{X}_{<t})\, d\boldsymbol{G}$.

This formulation motivates two coupled components. First, the temporal causal discovery component uses variational inference to learn an approximate graph posterior $q_\phi(\boldsymbol{G})$ for $p(\boldsymbol{G} \mid \boldsymbol{X}_{<t})$, capturing uncertainty in the directed dependencies among stocks. Second, the risk-aware FCM models $p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t}, \boldsymbol{G})$ by generating each stock's future return from its causal parents and a stochastic noise variable representing stock-specific risk. We train the two components jointly, letting the learned causal structure guide return prediction while observed returns supervise the causal discovery process.

4.2 Causal Discovery for Stock Relations

As illustrated in Figure 2(a), DISCO adopts a Bayesian approach to learning the temporal causal graph among stocks Gong et al. (2022). Given the historical market observations $\boldsymbol{X}_{<t}$, the graph structure is characterized by the posterior distribution $p(\boldsymbol{G} \mid \boldsymbol{X}_{<t})$. Exact posterior inference is intractable because the number of possible graph structures grows exponentially with the number of stocks and time lags. We therefore use variational inference Blei et al. (2017) to approximate this posterior with a tractable distribution $q_\phi(\boldsymbol{G})$, parameterized by ϕ Gong et al. (2022); Pamfil et al. (2020); Geffner et al. (2022). The graph prior and variational posterior are introduced below. After training, graphs sampled from $q_\phi(\boldsymbol{G})$ represent plausible directed relation structures in the market.

Graph Prior Following Gong et al. (2022); Geffner et al. (2022), we impose a sparse prior on the temporal causal graph. Since all edges point from past observations to the current return, the DAG constraint is naturally satisfied. We therefore encourage sparsity using the Frobenius norm: $p(\boldsymbol{G}) \propto \exp\left(-\sum_{l=1}^{L} \|G_l\|_F^2\right)$.

Variational Posterior As defined in Section 3.2, each variable $G_l^{ji} \in \{0,1\}$ indicates whether a directed causal edge exists from X_{t-l}^j to r_t^i. Since each candidate edge has only two possible states, present or absent, we model it as a Bernoulli random variable: $q_\phi\left(G_l^{ji}\right) \sim \text{Bernoulli}(\sigma_l^{ji})$, where $\sigma_l^{ji} \in [0,1]$ denotes the probability that the edge exists. Assuming that the candidate edges are conditionally independent under the variational approximation, the posterior distribution over the complete temporal causal graph is factorized as

$$q_\phi(\boldsymbol{G}) = \prod_{l=1}^{L}\prod_{j=1}^{D}\prod_{i=1}^{D} q_\phi\left(G_l^{ji}\right) = \prod_{l=1}^{L}\prod_{j=1}^{D}\prod_{i=1}^{D} \text{Bernoulli}\left(\sigma_l^{ji}\right). \tag{1}$$

This factorization allows all candidate edges to be sampled and evaluated in parallel, making graph inference scalable to large stock markets.

To parameterize the edge probabilities, we introduce two learnable tensors $\boldsymbol{U} = \{u_l^{ji}\}_{l=1:L}^{j,i=1:D}$ and $\boldsymbol{V} = \{v_l^{ji}\}_{l=1:L}^{j,i=1:D}$, both in $\mathbb{R}^{L\times D\times D}$. Here, u_l^{ji} and v_l^{ji} are the logits associated with the presence and absence, respectively, of the edge from X_{t-l}^j to r_t^i. The corresponding edge probability is obtained using a two-class softmax:

$$\sigma_l^{ji} = \exp\left(u_l^{ji}\right) / \left(\exp\left(u_l^{ji}\right) + \exp\left(v_l^{ji}\right)\right). \tag{2}$$

The variational parameters are therefore $\phi = \{\boldsymbol{U}, \boldsymbol{V}\}$. After training, $q_\phi(\boldsymbol{G})$ represents the distribution over the directed relation structure of the stock market.

Because the binary graph $\boldsymbol{G}$ indicates only whether each causal link exists, we additionally introduce a learnable causal-strength tensor $\widetilde{\boldsymbol{G}} = [\widetilde{G}_1, \widetilde{G}_2, \ldots, \widetilde{G}_L] \in \mathbb{R}^{L\times D\times D}$. Each element $\widetilde{G}_l^{ji}$ quantifies the strength of the corresponding link. The effective relation weights are given by $\boldsymbol{G} \odot \widetilde{\boldsymbol{G}}$ and are subsequently used by the FCM for return prediction. The learning objective is shown in Section 5.1.

4.3 FCM WITH DYNAMIC RISK MODELING

Given a sampled temporal causal graph $\boldsymbol{G}$, DISCO models the conditional return distribution $p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t}, \boldsymbol{G})$ using the additive-noise FCM introduced in Section 3.2. For each stock i, its future return is generated as

$$r_t^i = F_i\left(\mathbf{Pa}_{\boldsymbol{G}}^i(<t), z_t^i\right) = f_i\left(\mathbf{Pa}_{\boldsymbol{G}}^i(<t)\right) + z_t^i. \tag{3}$$

The deterministic function f_i aggregates market information from the causal parents of stock i, while the stochastic variable z_t^i represents stock-specific risk. The causal graph constrains the aggregator such that $\partial f_i / \partial X_{t-l}^j = 0$ whenever $X_{t-l}^j \notin \mathbf{Pa}_{\boldsymbol{G}}^i(<t)$.

Market Information Aggregator Given the historical market observations $\boldsymbol{X}_{<t} = \{X_{t-l}^i\}_{i=1,l=1}^{D,L} \in \mathbb{R}^{D\times L\times d}$, we apply a shared LSTM Hochreiter & Schmidhuber (1997) to the feature sequence of each stock. This produces price embeddings $\boldsymbol{E}_{<t} = \{E_{t-l}^i\}_{i=1,l=1}^{D,L} \in \mathbb{R}^{D\times L\times d_l}$, where d_l is the LSTM hidden dimension. These embeddings capture nonlinear temporal patterns in the historical price features. The information aggregated for stock i is determined jointly by the sampled binary graph $\boldsymbol{G}$ and the causal-strength tensor $\widetilde{\boldsymbol{G}}$. We parameterize the aggregator as

$$f_i\left(\mathbf{Pa}_{\boldsymbol{G}}^i(<t)\right) = \zeta_i\left(\sum_{l=1}^{L}\sum_{j=1}^{D} G_l^{ji}\widetilde{G}_l^{ji}\,\ell\left(E_{t-l}^j\right)\right), \tag{4}$$

where ℓ is an encoding network shared across stocks and time lags, and ζ_i is a stock-specific decoding network. The product $G_l^{ji}\widetilde{G}_l^{ji}$ ensures that information from X_{t-l}^j contributes to the prediction of r_t^i only when the corresponding causal edge is active.

Dynamic Risk Modeling The risk of a stock is commonly associated with the dispersion of its possible returns: a broader return distribution indicates greater uncertainty. We represent the

risk of stock i on day t using the stochastic component z_t^i, which is separated from the market-driven component f_i under the additive-noise formulation. Based on $r_t^i = f_i\left(\mathbf{Pa}_{\boldsymbol{G}}^i(<t)\right) + z_t^i$, the conditional return distribution could be induced directly by the noise distribution: applying the change-of-variables formula Khemakhem et al. (2021) gives

$$p\left(r_t^i \mid \boldsymbol{X}_{<t}, \boldsymbol{G}\right) = p_{z_i}\left(z_t^i\right)\left|\frac{\partial F_i}{\partial z_t^i}\right|^{-1} = p_{z_i}\left(z_t^i\right), \tag{5}$$

where $\left|\partial F_i / \partial z_t^i\right| = 1$ under the additive-noise formulation. Because the conditional return distribution inherits its shape from the noise distribution, the choice of p_{z_i} is important. Although Gaussian distributions are commonly adopted in classical mean–variance modeling Markowitz (1991), empirical stock returns often exhibit substantial asymmetry Zhao et al. (2020). We therefore model z_t^i using a skew-normal distribution, $p_{z_i}(z_t^i) = \mathrm{SN}(z_t^i; \lambda_i)$, where λ_i is a learnable stock-specific skewness parameter. A risk-aware return prediction is then generated by sampling z_t^i from this distribution and adding it to the aggregated market component, as illustrated in Figure 2(b)–(c).

5 Model Training & Inference

5.1 Training Objective

The training objective of DISCO combines probabilistic graph learning, return prediction, and ranking optimization. We jointly learn the variational graph posterior $q_\phi(\boldsymbol{G})$ and the FCM parameters by maximizing the evidence lower bound (ELBO) of the conditional log-likelihood $\log p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t})$.

Following Borovykh et al. (2017), we assume that the graph prior is independent of the input, such that $p(\boldsymbol{G} \mid \boldsymbol{X}_{<t}) = p(\boldsymbol{G})$. The ELBO is then written as

$$\begin{aligned}\mathcal{L}_{\mathrm{ELBO}} &= \mathbb{E}_{q_\phi(\boldsymbol{G})}\left[\log p(\boldsymbol{r}_t \mid \boldsymbol{X}_{<t}, \boldsymbol{G})\right] - \mathrm{KL}\left(q_\phi(\boldsymbol{G}) \,\|\, p(\boldsymbol{G})\right) \\ &= \mathbb{E}_{q_\phi(\boldsymbol{G})}\left[\sum_{i=1}^{D} \log p_{z_i}(z_t^i) + \log p(\boldsymbol{G}) - \log q_\phi(\boldsymbol{G})\right],\end{aligned} \tag{6}$$

where $z_t^i = r_t^i - f_i\left(\mathbf{Pa}_{\boldsymbol{G}}^i(<t)\right)$ is the residual noise associated with the ground-truth return r_t^i.

In addition to the probabilistic objective, we use a mean squared error loss to improve return prediction and a pairwise ranking loss to preserve the relative ordering of stocks:

$$\mathcal{L}_{\mathrm{MSE}} = \sum_{i=1}^{D}\left(\hat{r}_t^i - r_t^i\right)^2, \quad \mathcal{L}_{\mathrm{RANK}} = \sum_{i=1}^{D}\sum_{j=1}^{D}\max\left(0, -\left(\hat{r}_t^i - \hat{r}_t^j\right)\left(r_t^i - r_t^j\right)\right). \tag{7}$$

Finally, the overall training loss is

$$\mathcal{L} = -\mathcal{L}_{\mathrm{ELBO}} + \lambda_m \mathcal{L}_{\mathrm{MSE}} + \lambda_r \mathcal{L}_{\mathrm{RANK}}, \tag{8}$$

where λ_m and λ_r control the contributions of the prediction and ranking losses. We optimize the FCM parameters and the variational parameters ϕ jointly. Because the sampled graph is discrete, we use the Gumbel–Softmax reparameterization Maddison et al. (2016); Jang et al. (2016) to estimate gradients with respect to ϕ. The complete training procedure is provided in Algorithm 1 in Appendix D.

5.2 Multi-Situation Inference

Both the learned causal structure and stock-specific risk are uncertain at inference time. Inspired by n-shot sampling Yao et al. (2023), we introduce an optional multi-situation inference strategy, referred to as **DISCO with Multi-infer**, to account for these two sources of uncertainty. Uncertainty in market relations is represented by the graph posterior $q_\phi(\boldsymbol{G})$, while stock-specific uncertainty is represented by the risk distributions $\{p_{z_i}\}_{i=1}^D$.

Specifically, we first sample N causal graphs from $q_\phi(\boldsymbol{G})$, each representing a plausible directed relation structure. For each sampled graph, we draw M risk samples from p_{z_i} for every stock and generate the corresponding return predictions using the FCM. This process produces $N \times M$ predictions for each stock. We average these predictions to obtain the final return estimate and rank

stocks in descending order accordingly. When the graph posterior or risk distribution is concentrated, the corresponding samples exhibit less variation; more dispersed distributions yield more diverse situations. The resulting average therefore accounts for uncertainty in both market relations and stock-specific risk. The complete inference procedure is provided in Algorithm 2 in Appendix D.

5.3 Computational Complexity

We summarize the dominant computational cost of DISCO. Let B denote the number of trading-day samples in each training mini-batch, D the number of stocks, L the maximum time lag, and d_l the embedding dimension. During training, graph-based market aggregation dominates the computation, resulting in a per-step complexity of $\mathcal{O}(BD^2Ld_l)$. The training cost is therefore linear in B and L, and quadratic in D. During multi-situation inference, let N denote the number of sampled causal graphs and M the number of risks sampled for each graph. For every graph-risk combination, DISCO performs a complete FCM forward pass to generate a return prediction. Since each of the N graphs is evaluated under M risks, the inference procedure requires $N \times M$ forward passes, resulting in a complexity of $\mathcal{O}(NMD^2Ld_l)$. Setting $N = M = 1$ recovers the single-inference setting.

6 Experiments

6.1 Experimental Setup

We evaluate DISCO on three real-world daily stock datasets: NASDAQ with $1{,}026$ stocks, NYSE with $1{,}737$ stocks Feng et al. (2019b), and TSE with 95 stocks Wang et al. (2022). Each dataset is divided chronologically into training, validation, and test periods to prevent future information leakage. Following previous studies Cheng et al. (2023b), each method recommends the top-5 stocks on every test day. We evaluate performance using the Cumulative Investment Return Ratio (IRR), which measures overall profitability, and the Sharpe Ratio (SR), which measures risk-adjusted returns. We compare DISCO with 22 baselines covering stock movement prediction, price forecasting, reinforcement-learning-based trading, and stock recommendation. Detailed dataset statistics, metric definitions, baseline descriptions, and implementation settings are provided in Appendix B.

6.2 Main Results

Table 1 summarizes the main experimental results of two metrics on the three datasets. Overall, DISCO achieves the highest results against baselines across both metrics, which suggests improved risk-adjusted performance across different markets. As for profitability, DISCO shows a clear edge over baselines, which improves the *SR* of the SOTA baseline by $138.4\%, 120.7\%, 104.7\%$ with p-values less than 0.01 (one-tailed paired t-test under 5 runs), indicating statistical significance on all three datasets. Higher *SR* suggests better risk-adjusted returns, which is indicative of a more risk-tolerant trading strategy of DISCO. Greater degrees of improvements are observed on *NASDAQ* and *NYSE* datasets than *TSE*, which may indicate the advantage of DISCO in dealing with large candidate pools for stock recommendation. To evaluate the effectiveness of DISCO equipped with the multi-situation inference strategy, i.e. DISCO with Multi-infer, we also report its results and the standard deviations (Std, calculated across multiple inference results). By considering various market scenarios, DISCO with Multi-infer offers a more comprehensive assessment.

6.3 Ablation Study

To demonstrate the effectiveness of the proposed components, we compare DISCO with 7 variants: **DISCO w/o Causal**: the causal graph discovery process is replaced by an all-ones tensor to represent the relation; **DISCO w/o Risk**: remove the risk effect when calculating the return ratio; **DISCO with Normal Dist.**: replace the skew normal distribution with normal distribution; **DISCO with MLP / CNN / GRU / Transformer**: the LSTM module in FCM is replaced by a 3-layer MLP / 3-layer CNN / 3-layer GRU / 3-layer Transformer Encoder.

The ablation results are shown in Table 2, from which we can see that DISCO outperforms other variants in most markets. With discovered causal relations, DISCO could simulate the market causal information flow and make use of it, which benefits the accuracy of future return ratio assessment.

Table 1: Comparison results with baselines.

Method	NASDAQ		NYSE		TSE	
	SR	IRR	SR	IRR	SR	IRR
Adv-ALSTM	0.97	0.23	0.81	0.14	1.10	0.43
HGCluster	0.06	0.10	0.10	0.11	0.20	0.10
HATS	0.80	0.15	0.73	0.12	0.96	0.31
HMG-TF	0.83	0.19	0.75	0.13	1.05	0.33
LSTM-RGCN	0.75	0.13	0.70	0.10	0.90	0.28
HATR	0.92	0.31	0.76	0.14	0.98	0.36
CI-STHPAN	0.66	**2.10**	0.79	**2.14**	-	-
DA-RNN	0.71	0.14	0.66	0.13	0.86	0.25
DQN	0.93	0.20	0.72	0.12	1.08	0.31
SAE-LSTM	0.95	0.22	0.79	0.12	0.73	0.21
iRDPG	1.32	0.28	0.85	0.18	1.10	0.55
RAT	1.37	0.40	1.03	0.22	1.20	0.64
RSR-E	1.12	0.26	0.88	0.20	1.07	0.50
RSR-I	1.34	0.39	0.95	0.21	1.08	0.53
STHAN-SR	1.42	0.44	1.12	0.33	1.19	0.62
ALSP-TF	1.55	0.53	1.24	0.41	1.27	0.71
H-GAT	1.53	0.85	1.97	1.01	-	-
DHMoE	1.24	0.64	1.08	0.46	-	-
PTER	1.58	0.56	-	-	1.31	0.79
Enhancer	1.70	0.68	-	-	1.41	**0.90**
FinD§	2.32	1.03	2.23	0.89	-	-
SVAT-STHAN	3.10	0.59	2.61	0.38	-	-
DISCO	**7.39**†	**2.10**	**5.76**†	1.91	**2.68**†	0.88
DISCO with Multi-infer	**7.43**†	2.07	**6.23**†	2.05	**2.74**†	**0.90**
Std of Multi-infer	0.119	0.081	0.405	0.148	0.150	0.060

† indicates a significant improvement over the best baseline (one-tailed paired t-test over five runs).

Table 2: Ablation study results over 7 variants.

Method	NASDAQ		NYSE		TSE	
	SR	IRR	SR	IRR	SR	IRR
DISCO w/o Causal	5.90	1.72	5.01	1.53	1.95	0.61
DISCO w/o Risk	6.57	1.89	5.64	1.75	2.54	0.81
DISCO with Normal Dist.	7.36	2.02	5.83	1.83	2.08	0.64
DISCO with MLP	6.57	1.76	5.14	1.54	2.53	0.79
DISCO with CNN	5.91	1.71	5.00	1.50	1.92	0.69
DISCO with GRU	7.21	2.08	5.72	1.76	2.48	0.78
DISCO with Transformer	7.29	2.06	**5.85**	**1.93**	2.40	0.81
DISCO	**7.39**	**2.10**	5.76	1.91	**2.68**	**0.88**

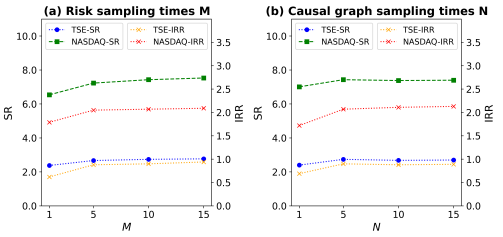

Figure 3: Sensitivity of DISCO to the Multi-infer sampling times. (a) The risk sampling times M vary with the causal-graph sampling times fixed at N=5. (b) The causal-graph sampling times N vary with M fixed at M=10.

By comparing the results of DISCO and DISCO w/o Risk, we find that taking the potential risk into account when assessing the future return ratio can indeed improve the profitability of DISCO. Moreover, the superior results of DISCO over DISCO with Normal Dist. indicate that using the skew normal distribution is more suitable for describing the distribution of return ratios and risks. For the Market Information Aggregator, modules with sequential representation capabilities (LSTM, GRU, and Transformer Encoder) offer clear advantages.

6.4 Model Analysis

The performance under different Top-K. We explore the impact of the performance of DISCO along with the number of selected top-rank stocks K. As shown in Figure 4, *SR* shows an upward trend as the selected top-rank stocks K increases, while *IRR* shows a downward trend. *SR* measures excess return per unit of risk. As the number of stocks in a portfolio increases, diversification tends to reduce the risk associated with specific stocks or sectors, leading to an upward tendency of *SR*. On the contrary, *IRR* generally reflects the overall profit without risk adjustment. Expanding the selected top-rank stock set means incorporating some with lower return ratios gradually, potentially reducing the overall return rate.

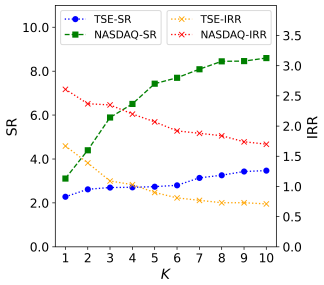

Figure 4: Top-K sensitivity.

Multi-situation Inference Strategy Study. We further examine the profitability of DISCO alongside sampling times for risk M and causal graph N. As shown in Figure 3(a), increasing risk sampling times leads to a rise in *SR* and stable *IRR* across all markets, indicating that sample risk multiple times effectively controls investment risks. Figure 3(b) demonstrates that profitability increases as the causal graph sampling times grow, highlighting the shifting market relations and underscoring the importance of considering various market situations. Table 3 reports the empirical *NASDAQ* latencies; even the heaviest configuration (M=1, N=15) costs only 0.046 s per sample on such a big market (1026 stocks). Per-dataset timing on all three markets is provided in Appendix C.

Table 3: Per-sample inference time on *NASDAQ*; cost is relative to Single-infer.

Config	Time (s)	Cost
Single-infer N=1, M=1	0.0044	1.00×
Multi-infer M=1, N=5	0.0163	3.71×
M=1, N=10	0.0308	7.04×
M=1, N=15	0.0460	10.49×
N=1, M=5	0.0139	3.17×
N=1, M=10	0.0255	5.83×
N=1, M=15	0.0375	8.55×
Final (N=5, M=10)	0.1228	27.91×

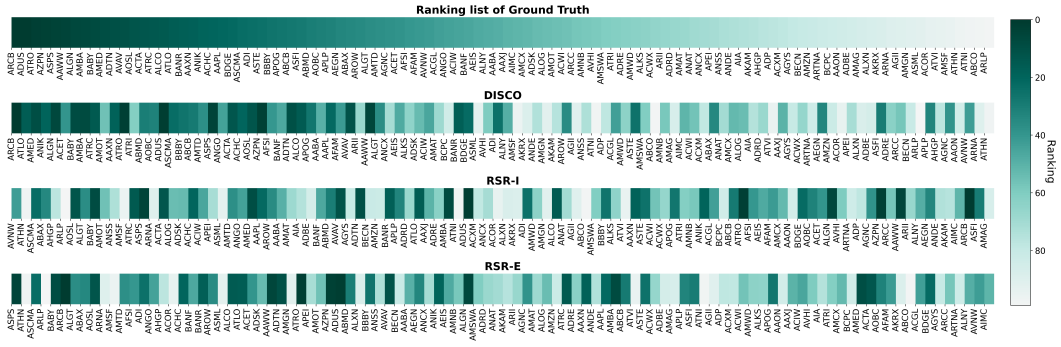

Figure 5: The case study of comparing the Top-100 recommendation ranking results generated from labels and methods on *NASDAQ* on the trading day *2017-01-03*. Darker colors represent higher rank.

Case Study of Stock Recommendation List. We present an example to illustrate the ranking list recommended by DISCO compared with the baseline methods RSR-E and RSR-I in Figure 5. We select the Top-100 stocks with the highest returns from the stock list of *NASDAQ* on the trading day *2017-01-03*. The darker colors represent higher-ranked stocks, and lighter colors indicate lower ranks. The figure shows that DISCO's ranking aligns more closely with the ground truth, exhibiting a smooth color transition from dark to light, while RSR-I and RSR-E deviate significantly. Additionally, we calculated Kendall's τ correlation coefficients Abdi (2007) to measure the ordinal association between the ground truth and each method's ranking. This coefficient ranges from -1 (indicating a negative ranking correlation) to $+1$ (indicating a positive ranking correlation), with 0 implying no correlation. DISCO achieved a coefficient of 0.0414, compared to 0.0008 for RSR-I and 0.005 for RSR-E, indicating DISCO's clear advantage.

Case Study of Causal Graph. Using the NASDAQ market as an example, Figure 6 presents the learned causal graph $\boldsymbol{G}$ multiplied by the causal strength graph $\tilde{\boldsymbol{G}}$ for 20 stocks, including 15 randomly selected ones and five tech giants: *Microsoft, Apple, NVIDIA, Adobe, Intel* in the technology industry. Darker colors represent stronger causal relationships. The cluster of darker shades in the upper left shows that stocks in the same industry influence each other more strongly. This causal graph provides explainable insights into stock relationships.

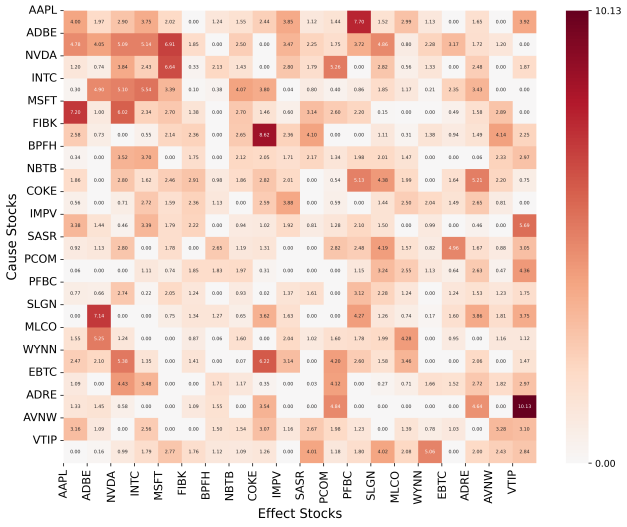

Figure 6: Stock relations visualization. Darker shades indicate stronger causal strength and vice versa.

7 Conclusion

This paper presents DISCO, a risk-aware stock recommendation framework that moves from correlation-based relation modeling to directed causal discovery. DISCO combines variational causal discovery with an additive-noise FCM to jointly model inter-stock dependencies, future returns, and stock-specific uncertainty. An optional multi-situation inference strategy further accounts for uncertainty in both the learned causal graph and the stochastic risk component. Experiments on NASDAQ, NYSE, and TSE show that DISCO improves the Sharpe ratio of the strongest baseline by more than 100% across all three markets while maintaining practical inference efficiency.

Reproducibility Statement

The implementation and hyperparameter settings are described in Appendix B, and the complete training and inference procedures are provided in Appendix D. Detailed empirical runtime measurements are reported in Appendix C. Anonymous source code is available at https://anonymous.4open.science/r/DISCO-D088.

AI Use Statement

Generative AI tools were used to assist with creating and revising scientific figures, suggesting the organization of the paper, and editing the manuscript for clarity and readability. The authors reviewed all AI-assisted content and verified that the revised figures accurately represent the underlying methods and results. Generative AI was not used to generate data, formulate scientific claims, develop the proposed method, conduct experiments, or interpret the results. The authors take full responsibility for the final content of the paper.

References

Hervé Abdi. The kendall rank correlation coefficient. *Encyclopedia of Measurement and Statistics. Sage, Thousand Oaks, CA*, pp. 508–510, 2007.

Wei Bao, Jun Yue, and Yulei Rao. A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PloS one*, 12(7):e0180944, 2017.

David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.

Anastasia Borovykh, Sander Bohte, and Cornelis W Oosterlee. Conditional time series forecasting with convolutional neural networks. *arXiv preprint arXiv:1703.04691*, 2017.

Salvatore M. Carta, Anselmo Ferreira, Alessandro Sebastian Podda, Diego Reforgiato Recupero, and Antonio Sanna. Multi-dqn: An ensemble of deep q-learning agents for stock market forecasting. *Expert Syst. Appl.*, 164:113820, 2021. URL https://api.semanticscholar.org/CorpusID:224963425.

Weijun Chen and Yanze Wang. Dhmoe: Diffusion generated hierarchical multi-granular expertise for stock prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 11490–11499, 2025.

Weijun Chen, Shun Li, Heyuan Wang, and Tengjiao Wang. Enhancer: A distribution-aware framework with temporal-relational meta-learning for stock prediction. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 250–261, 2025.

Jiezhu Cheng, Kaizhu Huang, and Zibin Zheng. Can perturbations help reduce investment risks? risk-aware stock recommendation via split variational adversarial training. *arXiv preprint arXiv:2304.11043*, 2023a.

Jiezhu Cheng, Kaizhu Huang, and Zibin Zheng. Can perturbations help reduce investment risks? risk-aware stock recommendation via split variational adversarial training, 2023b.

Rui Cheng and Qing Li. Subsequence-based graph routing network for capturing multiple risk propagation processes. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 3810–3816, 2022.

Qianggang Ding, Sifan Wu, Hao Sun, Jiadong Guo, and Jian Guo. Hierarchical multi-scale gaussian transformer for stock movement prediction. In *Ijcai*, pp. 4640–4646, 2020.

Fuli Feng, Huimin Chen, Xiangnan He, Ji Ding, Maosong Sun, and Tat-Seng Chua. Enhancing stock movement prediction with adversarial training. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pp. 5843–5849. International Joint Conferences on Artificial Intelligence Organization, 7 2019a. doi: 10.24963/ijcai.2019/810. URL https://doi.org/10.24963/ijcai.2019/810.

Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. Temporal relational ranking for stock prediction. *ACM Transactions on Information Systems (TOIS)*, 37(2): 1–30, 2019b.

Shibo Feng, Chen Xu, Yu Zuo, Guo Chen, Fan Lin, and Jianbing XiaHou. Relation-aware dynamic attributed graph attention network for stocks recommendation. *Pattern Recognition*, 121:108119, 2022.

Tomas Geffner, Javier Antoran, Adam Foster, Wenbo Gong, Chao Ma, Emre Kiciman, Amit Sharma, Angus Lamb, Martin Kukla, Nick Pawlowski, et al. Deep end-to-end causal inference. *arXiv preprint arXiv:2202.02195*, 2022.

Andreas Gerhardus and Jakob Runge. High-recall causal discovery for autocorrelated time series with latent confounders. *Advances in Neural Information Processing Systems*, 33:12615–12625, 2020.

Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.

Wenbo Gong, Joel Jennings, Cheng Zhang, and Nick Pawlowski. Rhino: Deep causal temporal relationship learning with history-dependent noise. *arXiv preprint arXiv:2210.14706*, 2022.

Olivier Goudet, Diviyan Kalainathan, Philippe Caillou, Isabelle Guyon, David Lopez-Paz, and Michele Sebag. Learning functional causal models with generative neural networks. *Explainable and interpretable models in computer vision and machine learning*, pp. 39–80, 2018.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.

Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems*, 21, 2008.

Yi-Ling Hsu, Yu-Che Tsai, and Cheng-Te Li. Fingat: Financial graph attention networks for recommending top-k k profitable stocks. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):469–481, 2021.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

Ilyes Khemakhem, Ricardo Monti, Robert Leech, and Aapo Hyvarinen. Causal autoregressive flows. In *International conference on artificial intelligence and statistics*, pp. 3520–3528. PMLR, 2021.

Raehyun Kim, Chan Ho So, Minbyul Jeong, Sanghoon Lee, Jinkyu Kim, and Jaewoo Kang. Hats: A hierarchical graph attention network for stock movement prediction. *arXiv preprint arXiv:1908.07999*, 2019.

Zengyu Lei, Caiming Zhang, Yunyang Xu, and Xuemei Li. Dr-gat: Dynamic routing graph attention network for stock recommendation. *Information Sciences*, 654:119833, 2024.

Wei Li, Ruihan Bao, Keiko Harimoto, Deli Chen, Jingjing Xu, and Qi Su. Modeling the stock relation with graph network for overnight stock movement prediction. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pp. 4541–4547, 2020. doi: 10.24963/ijcai.2020/626. Special Track on AI in FinTech.

Yang Liu, Qi Liu, Hongke Zhao, Zhen Pan, and Chuanren Liu. Adaptive quantitative trading: An imitative deep reinforcement learning approach. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 2128–2135, 2020.

Di Luo, Weiheng Liao, and Rui Yan. Lower risks, better choices: Stock correlation based portfolio selection in stock markets. In *Companion Proceedings of the ACM Web Conference 2023*, pp. 1–4, 2023.

Yongen Luo, Jicheng Hu, Xiaofeng Wei, Dongjian Fang, and Heng Shao. Stock trends prediction based on hypergraph modeling clustering algorithm. In *2014 IEEE International Conference on Progress in Informatics and Computing*, pp. 27–31, 2014. doi: 10.1109/PIC.2014.6972289.

Tao Ma and Ying Tan. Stock ranking with multi-task learning. *Expert Systems with Applications*, 199:116886, 2022.

Xiang Ma, Tianlong Zhao, Qiang Guo, Xuemei Li, and Caiming Zhang. Fuzzy hypergraph network for recommending top-k profitable stocks. *Information Sciences*, 613:239–255, 2022.

Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

Harry M Markowitz. Foundations of portfolio theory. *The journal of finance*, 46(2):469–477, 1991.

Jieyuan Mei, Jindong Tian, Ronghui Xu, Hanyue Wei, Chenjuan Guo, and Bin Yang. Find3: A dual 3d state space model with dynamic hypergraph for financial stock prediction. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pp. 2084–2094, 2025.

MYRON and S. Scholes. Crisis and risk management. *American Economic Review*, 2000.

Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. Dynotears: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics*, pp. 1595–1605. PMLR, 2020.

Yang Qiao, Yiping Xia, Xiang Li, Zheng Li, and Yan Ge. Higher-order graph attention network for stock selection with joint analysis. *arXiv preprint arXiv:2306.15526*, 2023.

Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison Cottrell. A dual-stage attention-based recurrent neural network for time series prediction. *arXiv preprint arXiv:1704.02971*, 2017.

Javier Sanz-Cruzado, Richard McCreadie, Nikolaos Droukas, Craig Macdonald, and Iadh Ounis. On transaction-based metrics as proxy for profitability of financial asset recommendations. In *Proceedings of the 3rd International Workshop on Personalization & Recommender Systems in Financial Services (FinRec 2022)*, 2022.

Ramit Sawhney, Shivam Agarwal, Arnav Wadhwa, Tyler Derr, and Rajiv Ratn Shah. Stock selection via spatiotemporal hypergraph attention network: A learning to rank approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 497–504, 2021. doi: 10.1609/aaai.v35i1.16127.

Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.

Peter Spirtes and Clark Glymour. An algorithm for fast recovery of sparse causal graphs. *Social science computer review*, 9(1):62–72, 1991.

Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.

Shuo Sun, Rundong Wang, and Bo An. Reinforcement learning for quantitative trading. *ACM Transactions on Intelligent Systems and Technology*, 14(3):1–29, 2023.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

Heyuan Wang, Shun Li, Tengjiao Wang, and Jiayi Zheng. Hierarchical adaptive temporal-relational modeling for stock trend prediction. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pp. 3691–3698, 2021. doi: 10.24963/ijcai.2021/508.

Heyuan Wang, Tengjiao Wang, Shun Li, Jiayi Zheng, Shijie Guan, and Wei Chen. Adaptive long-short pattern transformer for stock investment selection. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, pp. 3970–3977, 2022.

Heyuan Wang, Tengjiao Wang, Shun Li, Jiayi Zheng, Weijun Chen, and Wei Chen. Agree to disagree: Personalized temporal embedding and routing for stock forecast. *IEEE Transactions on Knowledge and Data Engineering*, 36(9):4398–4410, 2024.

Jingyuan Wang, Yang Zhang, Ke Tang, Junjie Wu, and Zhang Xiong. Alphastock: A buying-winners-and-selling-losers investment strategy using interpretable deep reinforcement attention networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1900–1908, 2019.

Hongjie Xia, Huijie Ao, Long Li, Yu Liu, Sen Liu, Guangnan Ye, and Hongfeng Chai. Ci-sthpan: Pre-trained attention network for stock selection with channel-independent spatio-temporal hypergraph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 9187–9195, 2024. doi: 10.1609/aaai.v38i8.28770.

Ke Xu, Yifan Zhang, Deheng Ye, Peilin Zhao, and Mingkui Tan. Relation-aware transformer for portfolio policy learning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pp. 4647–4653, 7 2020. Special Track on AI in FinTech.

Weiran Yao, Shelby Heinecke, Juan Carlos Niebles, Zhiwei Liu, Yihao Feng, Le Xue, Rithesh Murthy, Zeyuan Chen, Jianguo Zhang, Devansh Arpit, et al. Retroformer: Retrospective large language agents with policy gradient optimization. *arXiv preprint arXiv:2308.02151*, 2023.

Kun Zhang, Zhikun Wang, Jiji Zhang, and Bernhard Schölkopf. On estimation of functional causal models: general results and application to the post-nonlinear causal model. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(2):1–22, 2015.

Pengxiang Zhao, Shang Gao, and Nachuan Yang. Solving multi-objective portfolio optimization problem based on moea/d. In *2020 12th International Conference on Advanced Computational Intelligence (ICACI)*, pp. 30–37. IEEE, 2020.

Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems*, 31, 2018.

Zetao Zheng, Jie Shao, Jia Zhu, and Heng Tao Shen. Relational temporal graph convolutional networks for ranking-based stock prediction. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pp. 123–136. IEEE, 2023.

A Broader Impacts & Limitations

Limitations. DISCO is trained and evaluated on historical price data, and its predictive accuracy can degrade under *distribution drift*: macroeconomic regime changes, novel policy interventions, and abrupt market shocks (e.g., financial crises or unprecedented geopolitical events) can shift the underlying causal structure and risk distribution in ways that are not represented in the training window. The learned causal graph and risk parameters should therefore be periodically re-estimated, and out-of-distribution flags or re-training triggers should be considered before deploying the model in changing market conditions.

Broader Impacts. This work aims to advance machine learning methodologies for modeling temporal dependencies and uncertainty, with a particular focus on causal modeling and risk-aware stock recommendation. The contributions are primarily methodological in nature, and the experimental evaluation is conducted using historical, publicly available data. As with many machine learning studies, the proposed approach may have potential societal implications; however, these implications are consistent with those that are already well understood in the continued development of the field. We do not identify any new or specific ethical risks that warrant special emphasis in this paper. The proposed method is intended for research purposes. Any real-world deployment should be carefully evaluated in its specific context and should comply with applicable laws, regulations, and industry standards.

B Experimental Setup

Datasets. We evaluate DISCO on the following three real-world datasets. The detailed dataset statistics are shown in Table 4.

Table 4: Dataset statistics of the three datasets. # Dataset Partition denotes the number of time points in (Train, Validation, Test) split, respectively. Frequency denotes the sampling interval of time points.

	NASDAQ Feng et al. (2019b)	**NYSE** Feng et al. (2019b)	**TSE** Wang et al. (2022)
Markets	US	US	Japan
# Stocks	1026	1737	95
Train Period	01/2013-12/2015	01/2013-12/2015	11/2015-08/2018
Valid Period	01/2016-12/2016	01/2016-12/2016	08/2018-07/2019
Test Period	01/2017-12/2017	01/2017-12/2017	07/2019-08/2020
Frequency	Daily	Daily	Daily
# Dataset Partition (Train:Valid:Test)	756:252:237	756:252:237	693:231:235

- **NASDAQ** Feng et al. (2019b): NASDAQ is a US stock dataset which contains the historical daily price data (opening price, highest price, lowest price, closing price, trading volume) of 1026 stocks within the NASDAQ global and capital market.
- **NYSE** Feng et al. (2019b): NYSE includes the historical daily price data (opening price, highest price, lowest price, closing price, trading volume) of 1737 stocks in the New York Stock Exchange market.
- **TSE** Wang et al. (2022): TSE corresponds to the popular TOPIX-100 index, which includes the historical daily price data (opening price, highest price, lowest price, closing price, adjusted closing price, trading volume) of 95 stocks with the largest market capitalization in the Tokyo Stock Exchange Market.

Metrics. Following previous studies Cheng et al. (2023b), we adopt two metrics to evaluate the profitability of methods:

- *Cumulative Investment Return Ratio (IRR)* measures the sum of the daily return ratio of the recommended stocks on the test set. It is defined as:

$$IRR = \sum_t IRR^t = \sum_t \sum_{i \in K^t} r_t^i,$$

where $r_t^i = (p_t^i - p_{t-1}^i)/p_{t-1}^i$ denotes the return ratio of i-th stock on t-th trading day, and K^t denotes the recommended top-K stock set. We set $K = 5$, the same as the baselines.

- *Sharpe Ratio (SR)* is a risk-adjusted return metric describing the additional earnings an investor receives per unit of increase in risk.

$$\mathrm{SR} = \frac{\mathbb{E}\left[\mathrm{IRR}^t - R_f\right]}{\mathbb{S}\left[\mathrm{IRR}^t - R_f\right]},$$

where IRR^t is the investment return ratio on t-th trading day and R_f is the risk-free return. $\mathbb{E}$ and $\mathbb{S}$ denote the expectation and standard deviation operations.

Baselines. It is worth noting that financial market regulations prevent access to user transaction data, so we can't compare DISCO with those methods designed for classic recommendation tasks. Following the previous works Cheng et al. (2023a); Wang et al. (2022), we compare DISCO with 22 baselines, which can be divided into 4 categories: Adv-ALSTM Feng et al. (2019a), HGCluster Luo et al. (2014), HATS Kim et al. (2019), HMG-TF Ding et al. (2020), LSTM-RGCN Li et al. (2020), and HATR Wang et al. (2021) target the stock movement prediction task. DA-RNN Qin et al. (2017) and FinD^3 Mei et al. (2025) are methods that are designed for the stock price prediction task. DQN Carta et al. (2021), iRDPG Liu et al. (2020), RAT Xu et al. (2020) mainly employ reinforcement learning techniques to maximize the profit. Most importantly, there are 11 competitive baselines tailored for stock recommendation, including SAE-LSTM Bao et al. (2017), RSR-E Feng et al. (2019b), RSR-I Feng et al. (2019b), STHAN-SR Sawhney et al. (2021), ALSP-TF Wang et al. (2022), SVAT-STHAN Cheng et al. (2023b), H-GAT Qiao et al. (2023), CI-STHPAN Xia et al. (2024), PTER Wang et al. (2024), DHMoE Chen & Wang (2025) and Enhancer Chen et al. (2025).

Implementation Details. We train our model on 2 NVIDIA A100 80GB GPUs with Adam optimizer. The input feature size $d = 5$. To better explore the model's performance, we use grid search to decide on some training hyperparameters. We select the hidden size of price embedding from $[4, 8, 16]$ and get the best performance with size 8. The hidden size of MLPs ζ and ℓ in FCM is set to 332, while the LSTM module is set to 3 layers with hidden size $d_l = 332$. The learning rate is set as 0.0001 selected from $[0.01, 0.0001, 0.00001]$. The time lag L is set as 16 selected from $[4, 6, 8, 10, 12, 14, 16]$. The training step is set to 10^4. We set different batch sizes for three datasets, 4 for *NYSE*, 16 for *NASDAQ*, and 256 for *TSE*. The weights of loss terms are set as $\lambda_m = \lambda_r = 1$. For the multi-situation inference, the sampling times are set to $M = 10$ and $N = 5$ both selected from $[5, 10, 15]$. For the causal discovery parameters, the initial link likelihood $\boldsymbol{U}$ and $\boldsymbol{V}$ are set as an all-ones tensor. For the risk modeling, the initial mean and skewness of the Skew Normal distribution are set as 0, while its standard deviation is set as 0.1.

C Empirical Inference Efficiency

We evaluate the empirical inference efficiency of DISCO under different numbers of causal-graph and risk samples. Here, N denotes the number of causal graphs sampled from $q_\phi(\boldsymbol{G})$, and M denotes the number of risk realizations sampled for each graph. All measurements report GPU wall-clock time on a single NVIDIA A100 80GB GPU after five warm-up steps. Table 5 presents the average latency per trading-day sample and the total inference time over the complete test set.

Increasing either N or M raises the inference latency on the two larger markets, NASDAQ and NYSE, because more graph-risk combinations are evaluated. The increase is less pronounced on TSE, where the smaller stock universe allows the additional samples to be processed efficiently in parallel. For the same sampling count, increasing N is slightly more expensive than increasing M. For example, on NASDAQ, (N=10, M=1) requires 7.31 seconds for the complete test set, whereas (N=1, M=10) requires 6.05 seconds.

Under the default configuration (N=5, M=10), DISCO processes the complete test sets of NASDAQ, NYSE, and TSE in 29.10, 129.75, and 0.52 seconds, respectively. Even on NYSE, the largest dataset with 1,737 stocks, inference takes approximately 2.2 minutes for all 237 test days. These results indicate that multi-situation inference is practical for end-of-day stock ranking, where predictions are generated once per trading day. The reported measurements are specific to the stated hardware and implementation.

Table 5: Inference time of single-infer and multi-infer on the three datasets. Per-sample latency and total test-set time are reported in seconds.

Configuration	NASDAQ (1,026 stocks, 237 test days)		NYSE (1,737 stocks, 237 test days)		TSE (95 stocks, 235 test days)	
	Per sample	**Total**	**Per sample**	**Total**	**Per sample**	**Total**
Single-infer (N=1, M=1)	0.0044	1.04	0.0182	4.31	0.0017	0.41
Multi-infer (N=5, M=1)	0.0163	3.86	0.0639	15.14	0.0024	0.57
Multi-infer (N=10, M=1)	0.0308	7.31	0.1261	29.88	0.0031	0.72
Multi-infer (N=15, M=1)	0.0460	10.90	0.1878	44.51	0.0037	0.88
Multi-infer (N=1, M=5)	0.0139	3.29	0.0573	13.57	0.0019	0.44
Multi-infer (N=1, M=10)	0.0255	6.05	0.1110	26.30	0.0019	0.45
Multi-infer (N=1, M=15)	0.0375	8.89	0.1636	38.77	0.0019	0.44
Default (N=5, M=10)	**0.1228**	**29.10**	**0.5475**	**129.75**	**0.0022**	**0.52**

D ALGORITHMS

D.1 TRAINING ALGORITHM

The training process of DISCO is summarized in Algorithm 1.

D.2 MULTI-SITUATION INFERENCE ALGORITHM

The multi-situation inference process of DISCO is summarized in Algorithm 2.

Algorithm 1 The training process of DISCO.

Input: D stocks, time lag L, stock price feature dimension d. Training sample $\boldsymbol{X}_{<t}$ and $\boldsymbol{r}_t$. Batch size B, training step S, learning rate l_r, loss weights λ_m and λ_r.

Output: The approximate causal graph posterior $q_\phi(\boldsymbol{G})$ with parameter set ϕ, the market information aggregator f and the approximate risk distribution $p_z(z)$ with parameter set θ.

Initialize parameters ϕ and θ according to the Implementation Details section;
for training step = 1 to S **do**
  for each batch input including B samples **do**
    Sample the causal graph $\boldsymbol{G}$ from its posterior $q_\phi(\boldsymbol{G})$ as shown in Equation 2;
    Aggregate information based on the graph $\boldsymbol{G}$ and the input training samples by Equation 4;
    Sample the risk $\boldsymbol{z}$ from its distribution $p_{\boldsymbol{z}}(\boldsymbol{z})$;
    Predict ranking scores (return ratios) $\hat{\boldsymbol{r}}$ by Equation 3;
    Compute the ELBO term $\mathcal{L}_{ELBO}$ by Equation 6;
    Compute the MSE loss term $\mathcal{L}_{MSE}$ by Equation 7;
    Compute the ranking loss term $\mathcal{L}_{RANK}$ by Equation 7;
    Compute the overall loss term $\mathcal{L}$ by Equation 8;
    Update parameter set ϕ by $\phi \leftarrow \phi - l_r \cdot \nabla_\phi \mathcal{L}$;
    Update parameter set θ by $\theta \leftarrow \theta - l_r \cdot \nabla_\theta \mathcal{L}$;
  end for
end for
return the approximate causal graph posterior $q_\phi(\boldsymbol{G})$, the market information aggregator f and the approximate risk distribution $p_z(z)$.

Algorithm 2 Multi-situation inference of DISCO.

Input: Test dataset; graph posterior $q_\phi(\boldsymbol{G})$; market information aggregator f; risk distribution $p_{\boldsymbol{z}}(\boldsymbol{z})$; batch size B; graph-sampling count N; risk-sampling count M; recommendation size K.
Output: Top-K recommended stock indices for each trading day.

for each mini-batch containing B test samples **do**
  for $n = 1$ to N **do**
    Sample a causal graph $\boldsymbol{G}^{(n)} \sim q_\phi(\boldsymbol{G})$ using Equation 2.
    for $m = 1$ to M **do**
      Sample a risk realization $\boldsymbol{z}^{(n,m)} \sim p_{\boldsymbol{z}}(\boldsymbol{z})$.
      Aggregate market information from the test samples under $\boldsymbol{G}^{(n)}$ using Equation 4.
      Compute the return prediction $\hat{\boldsymbol{r}}^{(n,m)}$ using Equation 3.
    end for
  end for
  Compute $\bar{\boldsymbol{r}} = \frac{1}{NM} \sum_{n=1}^{N} \sum_{m=1}^{M} \hat{\boldsymbol{r}}^{(n,m)}$.
  Rank stocks in descending order of $\bar{\boldsymbol{r}}$.
  Record the Top-K stock indices for each trading day in the mini-batch.
end for
return the Top-K recommended stock indices for the test dataset.