Time Series Supplier Allocation via Deep Black-Litterman Model

Xinke Jiang[1*], **Wentao Zhang**[2*], **Yuchen Fang**[1*], **Xiaowei Gao**[3†], **Hao Chen**[4]
Haoyu Zhang[5], **Dingyi Zhuang**[6], **Jiayuan Luo**[7†]

[1] University of Electronic Science and Technology of China, Chengdu, China
[2] ShanghaiTech University, Shanghai, China
[3] University College London, London, United Kingdom
[4] University of Chinese Academy of Sciences, Beijing, China
[5] Zhongnan University of Economics and Law, Wuhan, China
[6] Massachusetts Institute of Technology, Cambridge, USA
[7] University of Macau, Macau, China
{thinkerjiangpku, wentaozh2001, fyclmiss}@gmail.com

Abstract

As a typical problem of Spatiotemporal Resource Management, Time Series Supplier Allocation (TSSA) poses a complex NP-hard challenge, aimed at refining future order dispatching strategies to satisfy the trade-off between demands and maximum supply. The Black-Litterman (BL) model, which comes from financial portfolio management, offers a new perspective for the TSSA by balancing expected returns against insufficient supply risks. However, the BL model is not only constrained by manually constructed perspective matrices and spatio-temporal market dynamics but also restricted by the absence of supervisory signals and unreliable supplier data. To solve these limitations, we introduce the pioneering **D**eep **B**lack-**L**itterman **M**odel (DBLM) for TSSA, which innovatively adapts the BL model from financial domain to supply chain context. Specifically, DBLM leverages Spatio-Temporal Graph Neural Networks (STGNNs) to capture spatio-temporal dependencies for automatically generating future perspective matrices. Moreover, a novel Spearman rank correlation is designed as our DBLM supervise signal to navigate complex risks and interactions of the supplier. Finally, DBLM further uses a masking mechanism to counteract the bias of unreliable data, thus improving precision and reliability. Extensive experiments on four datasets demonstrate significant improvements of DBLM on TSSA.

1 Introduction

As a typical problem of Spatiotemporal Resource Management, the goal of Time Series Supplier Allocation (TSSA) is to reduce discrepancies and boost efficiency by optimizing supplier capabilities to precisely match order quantities in the future (Wan, Rao, and Dong 2023; Alikhani, Torabi, and Altay 2019; Gören 2018). The TSSA optimization involves space and time thus can be seen as a typical spatiotemporal resource management problem but is a critical challenge in enhancing supply chain efficiency due to its NP-hard complexity (Nasiri et al. 2018; Kawtummachai and Van Hop 2005). (See Appendix A.1 for TSSA related work). Similar to the equilibrium-return dilemma in financial management (Markowitz 1952; Treynor 1962; Mossin 1966), TSSA focuses on maximizing profits while minimizing risks, highlighting the importance of supply-demand alignment in overall performance.

Inspired by the controllable profit variance assumption-based Markowitz model can optimize asset returns against portfolio risks (Fabozzi, Markowitz, and Gupta 2015; Steinbach 2001; Sharpe 1963), the Black-Litterman (BL) model integrates subjective investor insights as perspective matrices to enhance portfolio optimization, thus can generate more accurate portfolios that reflect expert predictions of expected returns across different investments (Black and Litterman 1991; Xing et al. 2018). However, the application of the BL model within supply chain management still faces distinct challenges. The manual creation of perspective matrices introduces uncertainty, struggling to capture real-time market conditions, and the non-linear complex dynamics of supplier-enterprise relationships (Kara, Ulucan, and Atici 2019; Islam, Amin, and Wardley 2021). A notable example of complex relationships is the difficulty in determining a supplier's maximum supply capabilities at specific moments (Alikhani, Torabi, and Altay 2019; Lee 2009), *i.e.* supplier A's capacity at time t_{i-1} in Figure 1.

From a practical standpoint, the quest for effective TSSA in supply chains marks a promising avenue for research (Tan, Kannan, and Handfield 1998; Chen, Lin, and Huang 2006). On the scientific front, the challenge extends beyond merely incrementally applying conventional BL models. It involves developing an effective model capable of generating comprehensive perspective matrices that accurately reflect supplier performance and dynamics among the supply-order gaps. Deep learning (DL) frameworks have emerged as leading solutions to capture non-linear correlations for analysis (Chai and Ngai 2020; LeCun, Bengio, and Hinton

[*]These authors contributed equally.
[†]Xiaowei Gao and Jiayuan Luo are corresponding authors. This work was done when Xinke was at UESTC.

2015). However, pioneering within the DL framework still presents three significant challenges in this specific domain:

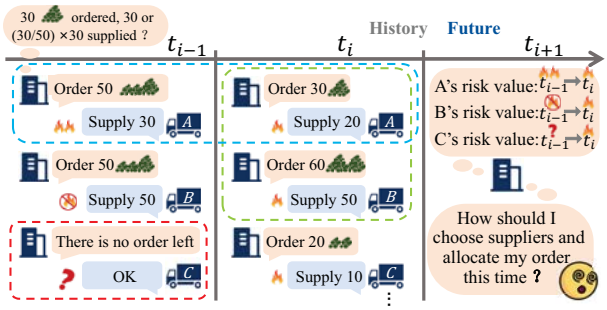

Figure 1: TSSA between enterprise and supplier. Fire denotes risk of shortfalls while a prohibition with it signifies no risk, and a question mark indicates unknown risks.

C1. Spatio-Temporal Dynamics. Supplier capacity exhibits inherent spatio-temporal dynamics crucial for future allocation. As depicted in Figure 1, at time t_i, the allocation of orders and supplies for suppliers A and B within the green rectangle shows contrasting trends: an increase for A whereas a decrease for B compared to their previous levels at time t_{i-1} respectively. These temporal trends indicate market-driven spatial competition and dynamics, which are essential for modeling future supplier capabilities. Yet, influenced by changing supplier and market inherence, these dynamics evolve over time, exemplified by A's diminished supply capability from t_{i-1} to t_i in the blue rectangle in Figure 1, highlighting the limitation of existing BL models to capture such changing spatio-temporal correlations, thus leading to inaccuracies in projecting future allocations.

C2. Lack of Supervisory Signals. Training DL models require robust supervisory signals to navigate gradient towards the optimal direction to objection (Ren, Guo, and Sutherland 2022). Yet, crafting these signals for comprehensive perspective matrices poses a major challenge, as accurate ground truths are hard and unavailable to establish, and supplementary data like supplier orders fall short in facilitating model training. Thus, our core focus is on creating appropriate supervisory signals for perspective matrix learning.

C3. Data Unreliability. Data unreliability, influenced by various factors, is a common issue in supply chain datasets, leading to biases in DL models, especially towards untraded or new suppliers with unassessed capabilities (Hao et al. 2022; Brockmann, Elson Kosasih, and Brintrup 2022). This is evident when supplier data, such as for untraded engaged suppliers, lacks historical engagement data, making their capabilities and risks uncertain. For example, as shown in the red rectangle in Figure 1, the absence of historical orders for supplier C obscures their supply potential and associated risks. Hence, integrating unreliable data for spatio-temporal analysis without a robust mechanism will introduce task-irrelevant biases, undermining the effectiveness of TSSA.

In light of the aforementioned technical limitations, we propose a novel **D**eep **B**lack-**L**itterman **M**odel (DBLM) to provide reasonable TSSA in this paper. Specifically, the BL model is first introduced to derive a historical analytic solution for the NP-hard supplier allocation task based on the profits of suppliers and the perspectives of enterprises. In DBLM paradigm, the perspective matrix of enterprises is data-driven instead of handcrafted, *i.e.*, the perspective matrix in our setting is learned by training a DL to mine historical supplier capacity and market dynamics. Inspired by recent spatio-temporal and sequential modeling methods (Fang et al. 2023; Li et al. 2022b; Jiang et al. 2023b; Liu et al. 2022), we combine the temporal and graph convolution network as the spatio-temporal graph neural networks (STGNNs) to learn the inherent spatio-temporal dynamics of historical supplier features and thus derive the perspective matrix with comprehensive knowledge to solve ***C1***. Therefore, the future solution of the supplier allocation, *i.e.*, the goal of our task, can be approximated through the spatio-temporal information-enhanced historical one. To introduce effective signals to address ***C2***, namely supervising the learning of future perspective matrices, we introduce a novel approach via a Spearman correlation coefficient-based ranking loss. This loss is designed around the negatively correlated weights and risks of suppliers, in line with that suppliers bearing lower risk should be awarded a greater proportion of orders. Moreover, we design a masking mechanism during calculating the ranking loss to mitigate the bias introduced by unreliable data to address ***C3***. Our contributions are as follows:

- To the best of our knowledge, our DBLM is the first initiative to integrate financial investment management strategies with supply chain demand challenges. We advanced the BL model to the TSSA task, pioneering the exploration of optimal solutions and modeling the perspective matrix through a spatio-temporal graph neural network.
- We introduce a novel masked ranking loss to guide DBLM training, which is implemented by Spearman rank correlation coefficient. Additionally, we implement a masking mechanism for unengaged suppliers within loss function to reduce bias from unreliable data.
- Our evaluation on four supplier allocation datasets demonstrates the superior performance of DBLM.

2 Preliminaries

Definition. (Order-Supply Mechanism Data) Consider an enterprise that consistently requires M volumes of raw materials, sourced from a set of suppliers denoted as $\mathcal{SU} = \{\mathcal{SU}_1, \mathcal{SU}_2, \cdots, \mathcal{SU}_N\}$, where N represents the total number of suppliers. This study examines the historical order demands and supply transactions between the enterprise and its suppliers over T periods. The order volumes are represented by $\mathcal{O} \in \mathbb{R}^{N \times T}$, where $\mathcal{O}_{it}$ specifies the order volume placed with supplier $\mathcal{SU}_i$ at time t. Similarly, supply volumes are denoted as $\mathcal{S} \in \mathbb{R}^{N \times T}$, with $\mathcal{S}_{it}$ indicating the supplied volume. Each unit volume sourced from supplier $\mathcal{SU}_i$ yields a constant return of $\mu_i (\mu_i \geq 0)$. However, supply volumes are inherently limited by the corresponding order volumes, i.e., $\forall \mathcal{SU}_i \in \mathcal{SU}, \mathcal{S}_{it} \leq \mathcal{O}_{it}$. This constraint leads to a "Shortage of Supply" (SoS) dilemma, which significantly impacts the continuity of production and sales operations. Consequently,

the enterprise is motivated to enhance its allocation strategies to meet the M volume requirement efficiently by perfectly matching the supply capacity of each individual supplier and also minimizing SoS risks, by analyzing historical order and supply patterns across the market.

Definition. (Black-Litter Model for Supplier Allocation) The supplier allocation challenge is conceptualized as managing a supplier portfolio, aimed at ensuring that the cumulative allocations across all suppliers in the portfolio meet the enterprise's total material requirement, M, while simultaneously minimizing the risk of SoS. The strategy specifies the allocation proportion or volume of orders $\mathcal{B}_i$ designated to each supplier $\mathcal{SU}_i$ within the portfolio for a given time period t. Importantly, these allocations are dynamic, and subject to adjustments in response to supplier performance, market conditions, and the changing needs of the enterprise. As such, the optimization task at any time t aims to:

$$\begin{cases} \max & \underbrace{\sum_{i=1}^{N} \mathcal{B}_{it}\mu_i}_{\text{Profit Item}} - \delta \underbrace{\sum_{i=1}^{N} \mathcal{B}_{it}^2 (\mathcal{O}_{it} - \mathcal{S}_{it})^{\kappa}}_{\text{Risk Item}} \\ \text{subject to} & \sum_{i=1}^{N} \mathcal{B}_{it} = M, i = \mathbb{N}^+. \quad \mathcal{B}_{it} \geq 0 \end{cases}, \quad (1)$$

where term $\sum_{i=1}^{N} \mathcal{B}it^2(\mathcal{O}it - \mathcal{S}it)^{\kappa}$ quantifies the risk, with κ being a positive integer exponent and δ as the coefficient balancing the profit and risk terms. This formula can be equivalently transformed, according to Appendix A.3, in alignment with the BL formula, as:

$$\begin{cases} \max & \underbrace{\mathcal{W}_t \mu}_{\text{Profit Item}} - \frac{\delta}{2} \underbrace{\mathcal{W}_t^T (\mathcal{O}_t - \mathcal{S}_t)^{\kappa} \mathcal{W}_t}_{\text{Risk Item}} \\ \text{subject to} & \sum_{i=1}^{N} w_{it} = 1, i = \mathbb{N}^+. w_{it} \in [0, 1] \end{cases}, \quad (2)$$

where $\mathcal{W}_t$ is the weight allocation matrix with entry $w_{it} = \mathcal{B}_{it}/M$ signifies the proportion of total material volume allocated to supplier $\mathcal{SU}_i$ at time t. Here we denote diagonal matrix $\Sigma_t = (\mathcal{O}_t - \mathcal{S}_t)^{\kappa}$, thus the optimist weight allocation matrix is solved as $\mathcal{W}_t^* = (\delta\Sigma)^{-1}\mu$, facilitating a strategic balance between maximizing profit, as demand here, and minimizing risk. However, notice that this optimization can only based on the existence of $\mathcal{O}, \mathcal{S}$, which is unachievable and not suitable for future projection. As a consequence, we introduce the DBLM to adjust the parameters μ (expected returns) and Σ (risk estimates) to better align with future conditions, effectively bridging the forecasting gap. This innovative approach allows for the recalibration of historical insights to forecast future allocation strategies more accurately, thereby enhancing the model's predictive reliability and strategic value in uncertain environments.

Definition. (Time Series Supplier Allocation) Given $\{\mathcal{O}, \mathcal{S}, M, \mu, N\}$ observed over the historical period from $t-p$ to t, the objective is to determine an optimal allocation strategy for the future time interval $t+1$ to $t+f$. The main goal of **TSSA** is to optimize the future allocation weights $\hat{\mathcal{W}}^* \in \mathbb{R}^{N \times f}$ of all N suppliers in future:

$$[\mathcal{S}_{t-p:t}, \mathcal{O}_{t-p:t}] \xrightarrow{\Theta} \hat{\mathcal{W}}^*_{t+1:t+f}, \quad (3)$$

where Θ is the learned parameters for DBLM. However, this decision-making process is complicated by the reliance on historical data pertaining only to orders $\mathcal{O}$ and supplies $\mathcal{S}$, alongside the presence of both subjective factors (e.g., the potential zeroing of order and supply data for non-engaged suppliers) and objective factors related to suppliers. Thus, achieving a globally optimal solution is impractical, rendering this an NP-hard problem as proved in (Meena and Sarmah 2013; Hu et al. 2018; Chauhan et al. 2023).

3 Method

In this section, we elaborate on our innovative DBLM as in Figure 2. Note that traditional models often struggle to incorporate the dynamic historical information of the market, leading to solutions that are optimized for past conditions but lack predictive power for future scenarios. To address this limitation, we introduce a data-driven subjective perspective from the enterprise on market dynamics. This approach allows for the adjustment of the profit term $\mu \rightarrow \hat{\mu}$ and the risk term $\Sigma \rightarrow \hat{\Sigma}$ in Section 3.3, facilitating the derivation of optimal solutions for future contexts. More specifically, we integrate traditional BL with advanced DL techniques by parameterizing the perspective matrix (denoted as $\mathcal{P}$) across both spatial and temporal dimensions through STGNN encoder. Based on BL principle, we obtain the resulting adjusted solution, denoted as $\mathcal{W}^*$, and then pass it through Predictor to form future allocation weight matrix $\hat{\mathcal{W}}^*$. To tackle the challenges (***C1: lack of supervisory signals*** and ***C2: data unreliability***) mentioned before, we construct a mask ranking loss. Notations are in Appendix A.2.

3.1 Data Preparation

Instead of learning spatio-temporal embeddings solely from unreliable data, we extract several crucial sequence-based features of suppliers as inputs in Section 3.1. Moreover, to capture market dynamics, we construct spatio-temporal interactions, formed by the relationships among suppliers to satisfy the spatial input requirements of STGNNs.

Feature Preparation. We construct eight quantitative supply indicators from order and supply volumes to serve as the initial features of the suppliers' capacity and stability in Appendix A.4. Then we concatenate the indicators for $\mathcal{SU}_i$ at time t into a comprehensive vector $f_{it} \in \mathbb{R}^8$ and denote $\mathcal{F}_t = \{f_{1t}, f_{2t}, \cdots, f_{Nt}\} \in \mathbb{R}^{N \times 8}$ for all suppliers and apply normalization to mitigate issues of scalar imbalance.

Dynamic Structure Construction. Furthermore, we construct a dynamic adjacency matrix to encapsulate the evolving interactions between suppliers. For each time point t, the Cosine similarity between suppliers $\mathcal{SU}_i$ and $\mathcal{SU}_j$ is computed as: $\text{sim}_t(i,j) = \frac{f_{it} \cdot f_{jt}}{\|f_{it}\| \cdot \|f_{jt}\|}$. Then derivation of the dynamic adjacency matrix $\mathcal{A}_t \in \mathbb{R}^{N \times N}$, is defined as: $\mathcal{A}_t(i,j) = \frac{\text{sim}_t(i,j)+1}{2}$, where $\mathcal{A}_t(i,j)$ signifies the connection strength between suppliers $\mathcal{SU}_i$ and $\mathcal{SU}_j$ at time t. To minimize the impact of noise connections, redundant edges with a weight below threshold 0.3 are pruned. Finally, normalized adjacency matrix, $\hat{\mathcal{A}}_t = \tilde{\mathcal{D}}_t^{-\frac{1}{2}} \mathcal{A}_t \tilde{\mathcal{D}}_t^{-\frac{1}{2}}$, with $\tilde{\mathcal{D}}_t$ as its degree matrix, serves as the feature propagation matrix (Jiang et al. 2024; Liu et al. 2021; Zhong et al. 2020; Yang et al. 2023; Jiang et al. 2023a; Li et al. 2022a).

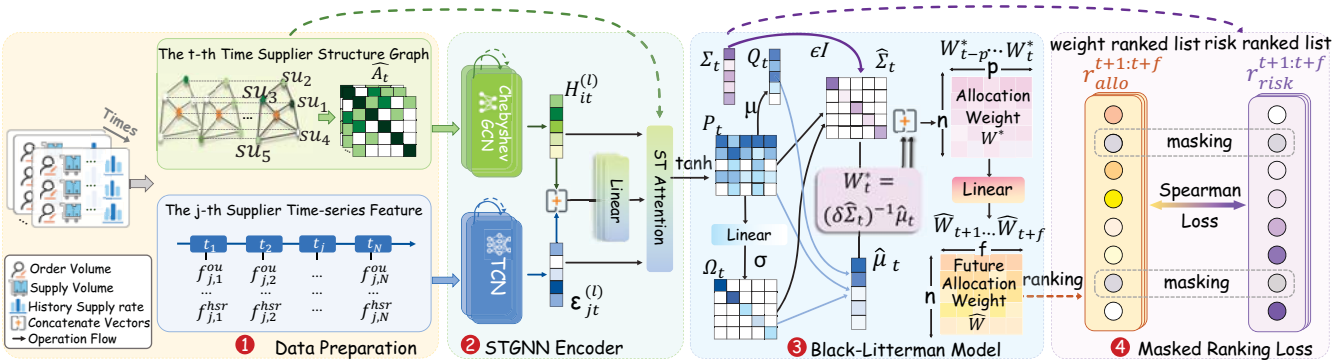

Figure 2: Overall framework of DBLM.

3.2 STGNN Encoder

In this section, we encode the prepared supplier sequence features $\{\mathcal{F}_{t-p}, \cdots, \mathcal{F}_t\}$ and dynamic propagation matrices $\widehat{A}_{t-p:t}$ to obtain representations in both spatial and temporal dimensions. We apply Spatial (Section 3.2) and Temporal Convolution (Section 3.2) to process these features and matrices, respectively. Such representations directly aid in the construction of the perspective matrix $\mathcal{P}$ (cf. Section 3.3).

Spatial Convolution Layer. The spatial block of DBLM is applied to enhance supplier correlation learning by leveraging global topological structure. In our implementation, we choose ChebGCN (Wu et al. 2021a) as Spatial Convolution:

$$\mathcal{H}_t^{(l)} = \sigma\left(\sum_{c=1}^{C} \mathbf{T}_c(\hat{\mathcal{A}}_t)\mathcal{H}_t^{(l-1)}\mathbf{W}_{\text{sp}}^{(l)}\right), \tag{4}$$

where $\mathcal{H}_t^{(l)}$ is the representation matrix in l-th layer at time t with $\mathcal{H}_t^{(0)} = \mathcal{F}_t$; the Chebyshev polynomial $\mathbf{T}_c(X) = 2X\mathbf{T}_{c-1}(X) - \mathbf{T}_{c-2}(X)$ is used to approximate the convolution operation, with boundary conditions $\mathbf{T}_0(X) = I$ and $\mathbf{T}_1(X) = X$; $\mathbf{W}_{\text{sp}}^{(l)}$ is the learnable spatial convolutional kernel in layer l, added to control how each node transforms the received information; $\sigma(\cdot)$ is the nonlinear activation function (e.g. ReLU).

Temporal Convolution Layer. The temporal block leverages the TCN to discern temporal trends from time-sequence features. Highlighted by (Wu et al. 2021b; Zhuang et al. 2022), TCNs offer significant improvements over RNNs due to the ability to process input sequences of variable lengths, enhancing adaptability across different temporal scales. The general idea of TCNs lies in utilizing a shared gated 1D convolution of width w_l in the l-th layer, facilitating the transmission of information from w_l adjacent time steps. As per Lea et al. (2017), each TCN layer H_l receives the signals from the previous layer H_{l-1} as:

$$\mathcal{E}_t^{(l)} = f\left(\mathbf{W}_{\text{te}}^{(l)} * \mathcal{E}_{t-1}^{(l-1)} + \mathbf{b}_{\text{te}}^{(l)}\right), \tag{5}$$

where $\mathcal{E}_t^{(l)}$ is the representation matrix in l-th layer at time t initialized with $\mathcal{E}_t^{(0)} = \mathcal{F}_t$; $\mathbf{W}_{\text{te}}^{(l)}$ is the convolution filter for the l-th layer, $*$ is the shared convolution operation, and $\mathbf{b}_{\text{te}}^{(l)}$ stands for the learnable bias.

3.3 Black-Litterman Predictor

After acquiring the spatial and temporal embeddings $\mathcal{H}_t^{(l)}, \mathcal{E}_t^{(l)}$ for each supplier, we need to base on the BL model to get the Perspective Matrix $\mathcal{P}$ and the Error Covariance Matrix Ω, which are critically important in our method and are detailed in the Appendix A.5. We employ a spatial-temporal fusion layer to derive $\mathcal{P}$ and Ω in Section 3.3. Thus, we can bridge the gap and reconcile the historical data with future expectations by adjusting the parameters μ and Σ with enriched spatio-temporal information. Consequently, we leverage $\mathcal{P}$ to solve and project the BL solution onto the future allocation matrix $\hat{\mathcal{W}}^*$, detailed in Section 3.3.

Spatio-Temporal Fusion Layer. Upon deriving the low-rank tensors $\mathcal{H}_t^{(l)}, \mathcal{E}_t^{(l)}$, we apply an attention mechanism (Fang et al. 2023) to fuse spatial and temporal representations, aiming to capture the complicated dynamics of supplier interactions. Achieved through the weighted attention-based low-tensor multiplication approach, $\mathcal{P}$ learns essential aspects of both spatial and temporal dimensions (Veličković et al. 2018; Zhang et al. 2024):

$$\begin{aligned} e_{ij,t} &= \text{LeakyReLU}\big(\vec{a}^{\mathsf{T}}\big[\mathbf{W}_{\text{attn}}[\mathcal{H}_{it}^{(l)}, \mathcal{E}_{jt}^{(l)}] + \mathbf{b}_{\text{attn}}\big]\big), \\ \alpha_{ij,t} &= \exp(e_{ij,t}) / \sum_{k=1}^{N} \exp(e_{ik,t}), \\ \mathcal{P}_t &= \tanh\Big(\sum_{\hat{\mathcal{A}}_t(i,j)>0} \alpha_{ij,t}\hat{\mathcal{A}}_t(i,j)[\mathcal{H}_{it}^{(l)} \times \mathcal{E}_{jt}^{(l)}]\Big), \end{aligned} \tag{6}$$

where $\alpha_{ij,t}$ is the spatio-temporal attention between $\mathcal{SU}_i$ and $\mathcal{SU}_j$ at time t. $\mathbf{W}_{\text{attn}}, \mathbf{b}_{\text{attn}}, \vec{a}$ are learnable matrices.

However, it is noted that $\mathcal{H}$ and $\mathcal{E}$ are not full-rank matrices, resulting in $\mathcal{P}$ also being non-full-rank (cf. Appendix A.11). Thus, due to data reliability, Σ is also non-full-rank due to zero-risk non-engaged suppliers, leading to traditional computation of Ω_t falls into non-full-rank, hindering the calculation of $(\mathcal{P}\tau\Sigma\mathcal{P}^T + \Omega)^{-1}$ in BL model, as detailed in Appendix A.12. To circumvent this dilemma, we adopt a calibration that incorporates the non-linear Sigmoid function and utilizes deep trainable matrices to refine and enhance the learning process of Ω, as detailed in Appendix A.12. After sigmoid calibration, the Ω_t is learned by:

$$\Omega_t = \text{diag}\big(\sigma(\mathbf{W}_{\text{om}}\mathcal{P}_t\Sigma_t\mathcal{P}_t^T + \mathbf{b}_{\text{om}})\big), \tag{7}$$

where $\mathbf{W}_{\text{om}}, \mathbf{b}_{\text{om}}$ are the learnable matrices. After applying $\text{diag}(\cdot)$ and $\sigma(\cdot)$, Ω_t transfers to a full-rank matrix. Additionally, to further enhance the capability to capture and integrate spatio-temporal dynamics, we employ a multi-head mechanism for learning $\mathcal{P}_t, \Omega_t$.

Black-Litterman Solver. The BL Model is enhanced by incorporating the perspective matrix $\mathcal{P}$ and error covariance matrix Ω from enterprise to adjust the return vector μ to $\hat{\mu}$ and risk vector Σ to $\widehat{\Sigma}$, as detailed in Appendix A.5. Given the equilibrium supplier profits Π normalized from μ via Appendix A.5, and perspective return vector $\mathcal{Q} = \mathcal{P} \times \mu + N(0, \Omega)$, we adjust the profit and risk components:

$$\begin{aligned} \widehat{\mu}_t &= \Pi_t + \tau\Sigma_t\mathcal{P}_t^T(\mathcal{P}_t\tau\Sigma_t\mathcal{P}_t^T + \Omega_t)^{-1}(\mathcal{Q}_t - \mathcal{P}_t\Pi_t), \\ \widehat{\Sigma}_t &= (1+\tau)(\Sigma_t + \epsilon I) - \tau(\Sigma_t + \epsilon I)\mathcal{P}_t^T \\ &\quad \times \left(\mathcal{P}_t\tau\Sigma_t\mathcal{P}_t^T + \Omega_t\right)^{-1}\mathcal{P}_t\tau\Sigma_t, \end{aligned} \tag{8}$$

where τ is the hyper-parameter that regulates the impact of enterprise's perspective $\mathcal{P}_t$ on adjustments. Moreover, we also make regularization calibration by adding ϵI to solve incalculaton of $\hat{\Sigma}^{-1}$ due to the non-full-rank problem caused by data unreliability (cf. Appendix A.13). By integrating $\widehat{\mu}_t, \widehat{\Sigma}_t$ into the profit and risk components, the solution:

$$\mathcal{W}_t^* = (\delta\widehat{\Sigma}_t)^{-1}\widehat{\mu}_t, \tag{9}$$

where $\mathcal{W}_t^*$ denotes the history optimal allocation strategy, ensuring $\sum \mathcal{W}_t^* = 1$. However, it is still struggling to adapt this historically optimal BL solution to align with future projections. Since future optimal allocation exhibits intricate relationship gaps with $\mathcal{W}_t^*$. Consequently, for mapping the optimal allocation into future intervals (Chen et al. 2023; Zhang et al. 2022), a non-linear transformation is applied on the combination of previous allocations $\{\mathcal{W}_{t-p}^*, \cdots, \mathcal{W}_t^*\}$:

$$\hat{\mathcal{W}}_{t+1:t+f}^* = \text{Softmax}(\mathbf{W}_{\text{out}}\mathcal{W}_{t-p:t}^* + \mathbf{b}_{\text{out}}), \tag{10}$$

where $\mathbf{W}_{\text{out}}, \mathbf{b}_{\text{out}}$ are learnable matrices. $\text{Softmax}(\cdot)$ is applied time-wisely, to ensure the sum of allocation weights normalized to 1.

3.4 Loss Function

Due to the challenges posed by ***C1: lack of supervisory signals*** and ***C3: data unreliability***, traditional loss functions like Mean Squared Error (MSE) (Wang and Bovik 2009) and Cross-Entropy (CE) (De Boer et al. 2005) tend to underperform. Consequently, we introduce a novel mask ranking loss crafted-carefully to address the two challenges. We denote the risk list $(O_t - S_t)^\kappa$ as risk ranked list $\mathbf{r}_{\text{risk}}^t$ in ascending order and predicted allocation weight $\hat{\mathcal{W}}_t^*$ as allocation ranked list $\mathbf{r}_{\text{allo}}^t$ in descending order. The two lists are converted to sequences of ranks, where r_i^t indicates the rank of the i-th item at time t. Then, the two ranks are leveraged to compute the negative part of the Spearman correlation coefficient (Mitra and Zhang 2014) as supervised loss $\mathcal{L}$. By minimizing $\mathcal{L}$, we can optimize the monotonic trend between two ranked sequences to address ***C1***:

$$\min_{\Theta}\mathcal{L} = \sum_{t=0}^{T_{train}}\sum_{j=t}^{t+f}\frac{6\sum_{i=1}^{N-|\mathcal{S}_j|}(r_{i,\text{risk}}^j - r_{i,\text{allo}}^j)^2}{(N-|\mathcal{S}_j|)\left((N-|\mathcal{S}_j|)^2-1\right)} + \eta||\Theta||^2, \tag{11}$$

where $N - |S_j|$ indicates we only calculate the gradients where $\mathcal{S}_{ij} \neq 0$ to solve ***C3***. Weight decay regularization $||\Theta||^2$ weighted by η is applied to mitigate overfitting. In the implementation, we adopt a masking mechanism, and due to the loss of gradient information after ranking, we utilize soft rank as suggested by (Blondel et al. 2021). Training complexity and algorithm in Appendix A.6 A.7.

4 Experiments

In this section, we explore the efficacy of DBLM in optimizing allocation task for time-series suppliers on four benchmark datasets [1]. We aim to address the following research questions: (1)**RQ1:** How does DBLM compare with SOTA methods in optimizing allocation SoS risk for time-series suppliers? (2)**RQ2:** What contributions do the key components of DBLM make to improving supplier allocation outcomes? (3)**RQ3:** Can the learned perspective matrix through STGNNs enhance the optimization of allocation risk? (4)**RQ4:** How does DBLM's performance vary with adjustments to the risk coefficient δ and reweight coefficient η?

In addition, we also conduct a robustness analysis in Appendix A.8 and hyper-parameter study in Appendix A.9.

4.1 Experiment Setup

Datasets. We evaluate DBLM on four supply chain datasets to optimize Time Series Supplier Allocation (TSSA), namely *MCM* and *SZ*, to ensure comprehensive validation. The *MCM* dataset comprises supply and order data from 401 suppliers over 240 weeks, and the *SZ* dataset includes data from 218 suppliers across 2 years (731 days).

Compared Methods. We benchmark DBLM against 14 methods across the *MCM-TSSA* and *SZ-TSSA* datasets:

- **Basic Algorithms: HA** (Guo et al. 2019) and **MC** (Puvska and Stojanovic 2022).
- **Decision-making Algorithms: Greedy** (Zhao et al. 2023), **DP** (Kuroiwa and Beck 2023), **Fuzzy-AHP** (Rezaei et al. 2020), **Fuzzy-TOPSIS** (Hasan et al. 2020) and **Markowitz** (Way et al. 2018).
- **Machine Learning Methods: DT** (Bowser-Chao and Dzialo 1993), **Lasso** (Tibshirani 1996), **MLP** (Rosenblatt 1963), **ECM** (Xing et al. 2018), **SGOMSM** (Hui et al. 2023) and **AGA** (Li et al. 2021).

Basic and decision-making algorithms aim at identifying historically risk-optimal portfolios, whereas the machine learning methods, including DBLM, utilize predictive analytics to assemble portfolios optimized for future allocation risks. Implementation details are in Appendix A.14.

Evaluation Metrics. Given the challenges highlighted in ***C3*** regarding the unreliability of predicted ground-truth allocation data, traditional evaluation metrics such as MAE and RMSE are deemed less suitable for our context. Therefore, we opt for two alternative evaluation metrics, **Hit Ratio@K (HR@K)** and **Mask Risk Expect (MRE)**, to more accurately and fairly evaluate the model's performance. Details

[1] Source codes and appendix are openly accessible at https://github.com/QiuFengqing/DBLM.

Method	Dataset	MCM-TSSA				SZ-TSSA			
	Metric	HR@10	HR@20	HR@50	MRE	HR@10	HR@20	HR@50	MRE
Baselines	HA	0.045±0.087	0.125±0.058	0.268±0.049	0.968±0.092	0.039±0.086	0.104±0.117	0.230±0.099	0.929±0.087
	MC	0.053±0.096	0.148±0.072	0.276±0.087	0.924±0.053	0.059±0.147	0.141±0.152	0.245±0.104	0.859±0.057
	Greedy	0.078±0.050	0.166±0.061	0.307±0.044	0.902±0.108	0.072±0.088	0.154±0.120	0.349±0.109	0.995±0.149
	DP	0.075±0.082	0.155±0.070	0.303±0.053	0.930±0.124	0.069±0.075	0.137±0.096	0.346±0.142	0.942±0.155
	Fuzzy-AHP	0.204±0.197	0.241±0.132	0.311±0.155	0.897±0.162	0.169±0.098	0.217±0.133	0.306±0.129	0.742±0.140
	Fuzzy-TOPSIS	0.104±0.128	0.187±0.140	0.233±0.165	0.887±0.143	0.095±0.087	0.127±0.094	0.149±0.138	0.939±0.143
	Markowitz	0.139±0.170	0.227±0.158	0.309±0.106	0.997±0.191	0.118±0.149	0.154±0.110	0.289±0.128	0.844±0.185
	DT	0.040±0.492	0.098±0.524	0.204±0.460	0.974±0.680	0.038±0.612	0.106±0.598	0.206±0.720	0.977±0.749
	Lasso	0.066±0.544	0.137±0.670	0.296±0.399	0.872±0.721	0.061±0.482	0.161±0.670	0.350±0.648	0.736±0.725
	MLP	0.199±0.344	0.245±0.287	0.331±0.225	0.973±0.339	0.182±0.291	0.246±0.348	0.382±0.306	0.556±0.320
	ECM	0.272±0.282	0.289±0.299	0.348±0.310	0.641±0.407	0.253±0.238	0.290±0.288	0.412±0.271	0.493±0.377
	SGOMSM	0.263±0.397	0.311±0.403	0.327±0.454	0.844±0.429	0.204±0.140	0.282±0.198	0.369±0.245	0.671±0.298
	AGA	0.158±0.237	0.206±0.228	0.310±0.296	0.772±0.357	0.180±0.205	0.242±0.167	0.374±0.152	0.629±0.261
Ours	DBLM	**0.403**±0.284	**0.449**±0.293	**0.487**±0.356	**0.518**±0.292	**0.481**±0.158	**0.543**±0.187	**0.662**±0.182	**0.327**±0.323
Ablation	DBLM(w/o BL)	0.154±0.488	0.238±0.462	0.347±0.529	0.820±0.442	0.112±0.658	0.148±0.495	0.325±0.431	0.729±0.480
	DBLM (w/o STGNN)	0.306±0.280	0.348±0.305	0.377±0.340	0.852±0.319	0.274±0.144	0.309±0.195	0.420±0.170	0.438±0.266
	DBLM (w/o TCN)	0.314±0.277	0.370±0.284	0.393±0.342	0.648±0.320	0.293±0.109	0.341±0.132	0.448±0.240	0.361±0.258
	DBLM (w/o DGCN)	0.323±0.211	0.419±0.277	0.431±0.330	0.719±0.376	0.340±0.172	0.442±0.249	0.473±0.150	0.377±0.342
	DBLM (w/o Fusion)	0.379±0.299	0.420±0.311	0.453±0.328	0.588±0.347	0.364±0.455	0.425±0.298	0.588±0.211	0.367±0.243
	DBLM (w/o Mask)	0.376±0.280	0.390±0.246	0.426±0.359	0.626±0.341	0.349±0.240	0.426±0.328	0.539±0.331	0.427±0.397
	DBLM (w/o Rank Loss)	0.290±0.279	0.317±0.194	0.335±0.243	0.692±0.287	0.307±0.198	0.373±0.276	0.501±0.453	0.486±0.493

Table 1: Performance comparison on *MCM* and *SZ*. The best is in red, and the second is in green in baselines, blue in ablations.

of evaluation matrics can be referred to Appendix A.10. In our implementation, we select HR@10, HR@20, HR@50, and MRE as the key metrics for evaluation.

Reproducibility. To ensure reproducibility, we optimize the parameters of baseline models using the Adam Optimizer with L_2 regularization and a dropout rate of 0.2. The sequence length for both input and forecasting is set to $p = f = 4$. For the DBLM model and baselines incorporating GNNs and TCNs, we utilize three layers with 150 hidden units each. For the DBLM model specifically, we set $\tau = 3, \delta = 0.6, \eta = 1e^{-4}$, and $\kappa = 2$, with the number of attention heads at 3 and a soft rank regularization strength of 0.5. An early-stopping strategy with a patience of 10 epochs is employed to mitigate overfitting. The dataset is divided into 70% training, 10% validation, and 20% testing portions. Implementations are done using the PyTorch 1.9.0 in Python 3.8 on NVIDIA Tesla V100 GPU. ECM, SGOMSM, and AGA are implemented using author-provided source codes, and other baselines are all implemented by ourselves.

4.2 Result Analysis

Main Results (RQ1) We conduct experiments and report performance on *MCM* and *SZ* datasets in Table 3. From the reported HR@K and MRE, we find the following observations: **(1) Comparison with Traditional Methods.** Decision-making algorithms (Greedy, DP, Fuzzy-AHP, Fuzzy-TOPSIS, Markowitz) demonstrate superior performance over basic approaches (HA, MC) in terms of HR@K and MRE. This indicates that utilizing historical data can enhance future risk allocation strategies to some extent. However, these algorithms are considerably less effective than machine learning models, underscoring the critical importance of leveraging predictive analytics for addressing the TSSA challenge ***C1***. This highlights the limitations of decision-making and basic methods, which rely solely on historical data. While DT, Lasso, MLP, and AGA attempt to capture temporal dependencies, their failure to adequately address the complex spatio-temporal interplay results in suboptimal outcomes, which can not reveal the variance in the supplier market. ECM and SGOMSM, optimized for allocation tasks and utilizing STGNNs, significantly outperform other baselines but falter in managing data unreliability, treating unknown or unordered suppliers as zero or no risk, which severely includes bias and undermines their efficacy in pinpointing high-risk entities. In contrast, DBLM leverages a Mask Rank Loss to deal with data unreliability and incorporates a perspective matrix that spans both spatial and temporal dimensions, offering a more nuanced understanding of market dynamics, and surpassing all baselines. **(2) Consistent Performance Superiority.** DBLM consistently outperforms all baselines, with improvements ranging from 19.2% to 48.2% on the *MCM* dataset and 33.7% to 90.1% on the *SZ* dataset. Notably, the model's performance excels at higher HR@K values, especially for larger K values, such as HR@50. For instance, DBLM secures an impressive hit ratio of 48.7% on *MCM* and 66.2% on *SZ*, emphatically demonstrating the superiority of utilizing STGNNs to address the TSSA issue. This capability significantly empowers decision-makers to sidestep high-risk suppliers, offering critical insights from both market and temporal perspectives.

Ablation Study (RQ2) We perform ablation studies to evaluate the impact of each component within DBLM, summarized in Table 1, with seven variants as detailed in Appendix A.10. It reveals that each component makes a positive contribution to DBLM's overall performance. Notably, the removal of the BL results in the most pronounced performance degradation. This underlines BL's vital role in modeling the perspective matrix within DBLM. Furthermore, the

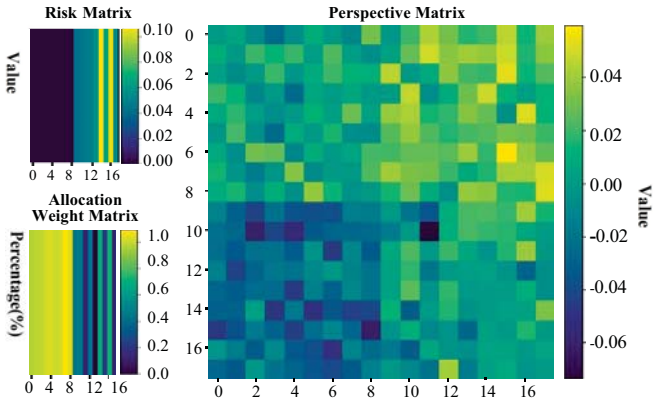

Figure 3: Risk Matrix (**Left**) composed of the top and bottom 9 suppliers in ascending sort, with corresponding Allocation Weight (**Right**) and Perspective Matrix (**Middle**).

absence of STGNN, TCN, or DGCN modules illustrates the necessity of capturing spatio-temporal dependencies to effectively tackle TSSA challenge, in line with ***C1***. Notably, the removal of DGCN has a slightly less detrimental effect than the absence of TCN module, emphasizing critical importance of temporal dynamics in the TSSA task. Moreover, the variant without Rank Loss exhibits the second-largest decrease in performance metrics such as HR@K and MRE, following the variant without BL. This indicates the lack of supervisory signals, ***C2***, significantly impairs the model's effectiveness, highlighting critical role of adopting appropriate loss functions. This rationale similarly applies to w/o Mask, providing a profound understanding of the contributions of Mask Mechanism to address ***C3***.

Case Study (RQ3) We visualize allocation weights $\hat{\mathcal{W}}$ and explain the learned perspective matrix $\mathcal{P}$ from the initial time horizon of the test set. Indeed, we concentrate on distinguishing between the top and bottom 9 suppliers ranked by ground-truth risk, after initially filtering our unreliable data. We categorize suppliers numbered 0-8 as those embodying the lowest risk (top 9) and suppliers numbered 9-17 as those representing the highest risk (bottom 9), arranging them in ascending order from the lowest to highest risk values. Subsequently, we correlate each supplier with their respective allocation weight referencing their risk rank. The illustrations in Figure 3, its left and right subgraphs—demonstrate that the allocation weight percentages with the lowest risk, conspicuously exceed those for the higher-risk suppliers, with values nearing 1%. This observation is in line with our theoretical assertion that suppliers bearing lower risk should be awarded a greater proportion of orders. Moreover, we analyze the relative perspective from the decision-maker of the top and bottom 9 suppliers based on risk ranking, which is in the middle of Figure 3. Within this matrix, lighter shades of $\mathcal{P}_{ij}$ indicate supplier $\mathcal{SU}_i$ enjoys higher competitiveness compared to $\mathcal{SU}_j$, whereas darker shades signal lower competitiveness to $\mathcal{SU}_j$. Observations from the matrix reveal that shading near upper-left and lower-right corners, tends towards values close to 0, indicating that the differences within each group (top or bottom 9 suppliers) are considerably less pronounced than those between two groups. Whereas, lighter shades are predominantly in upper-left quadrant and darker shades in lower-right, suggesting that top 9 suppliers, by their nature, are superior and more competitive than the bottom 9, which coincidentally corresponds to high-risk and low-risk ones. The insightful comparison within $\mathcal{P}$ provides empirical validation for our theoretical proposition that suppliers with lower risk are inherently more competitive in the market.

Dataset	SLD-TSSA			BSS-TSSA		
Model	HR@10	HR@20	HR@50	HR@10	HR@20	HR@50
Greedy	0.154	0.307	0.427	0.022	0.049	0.104
AGA	0.145	0.292	0.438	0.040	0.056	0.110
MLP	0.130	0.284	0.411	0.018	0.049	0.104
SGOMSM	0.127	0.336	0.435	0.038	0.056	0.094
ECM	0.127	0.278	0.401	0.033	0.057	0.108
DBLM	**0.315**	**0.464**	**0.553**	**0.043**	**0.065**	**0.112**

Table 2: Performance comparison on Traffic Management Datasets *SLD* and *BSS*.

Generalization Study in Traffic Domain To demonstrate the generalizability and effectiveness of our DBLM, we extend our experimentation to the traffic management domain. Specifically, for any given spatial segment (zone or station), the actual inflow and outflow within any specific time interval remains stochastic and unpredictable, which establishes a strong analogy to the TSSA problem. We model the vehicle inflow during specified time intervals as the enterprise's order quantity, while the vehicle outflow represents the supplier's supply quantity. We use the SLD dataset (Zhuang et al. 2022) (67 zones, 3-month duration, with 5-minute time intervals) and the BSS dataset (Gao, Chen, and Haworth 2023) (797 stations, 10,207,268 trips, and over 3 years of data with 15-minute time intervals) for the generalization experiment. From the above Table 2, DBLM demonstrates superior performance on the two datasets, showcasing strong generalization capabilities across different spatial-temporal resolutions and scenarios.

5 Conclusions

DBLM makes a significant advancement in spatiotemporal resource management of time series supplier allocation. By integrating BL model with the innovative perspective matrix, we successfully transform this complex issue into a manageable risk optimization problem, capturing both temporal and spatial information in market dynamics. The implementation of STGNNs of supplier relationship has notably streamlined the process, offering a more profound understanding of supplier inter-dependencies. The introduction of masked ranking mechanism, designed to address the lack of supervisory signals and unreliable issues, has further refined our approach, leading to more efficient supplier ranking and decision-making processes. Tested on four datasets, DBLM demonstrates superior performance over the traditional models and sets a new benchmark.

References

Alikhani, R.; Torabi, S.; and Altay, N. 2019. Strategic supplier selection under sustainability and risk criteria. *Int J Prod Econ*.

Black, F.; and Litterman, R. B. 1991. Asset Allocation. *J. Fixed Income*.

Blondel, M.; Teboul, O.; Berthet, Q.; and Djolonga, J. 2021. Fast Differentiable Sorting and Ranking. In *ICML*.

Bowser-Chao, D.; and Dzialo, D. L. 1993. Comparison of the use of binary decision trees and neural networks in top-quark detection. *Phys. Rev. D*.

Brockmann, N.; Elson Kosasih, E.; and Brintrup, A. 2022. Supply Chain Link Prediction on Uncertain Knowledge Graph. In *SIGKDD*.

Chai, J.; and Ngai, E. W. 2020. Decision-making techniques in supplier selection: Recent accomplishments and what lies ahead. *Expert Syst. Appl*.

Chauhan, V. K.; et al. 2023. Real-time large-scale supplier order assignments across two-tiers of a supply chain with penalty and dual-sourcing. *CIE*.

Chen, C.-T.; Lin, C.-T.; and Huang, S.-F. 2006. A fuzzy approach for supplier evaluation and selection in supply chain management. *Int J Prod Econ*.

Chen, S.-A.; Li, C.-L.; Yoder, N.; Arik, S. O.; and Pfister, T. 2023. TSMixer: An All-MLP Architecture for Time Series Forecasting. arXiv:2303.06053.

De Boer, P.-T.; Kroese, D. P.; Mannor, S.; and Rubinstein, R. Y. 2005. A tutorial on the cross-entropy method. *ANN OPER RES*.

Fabozzi, F. J.; Markowitz, H. M.; and Gupta, F. 2015. Portfolio Selection.

Fang, Y.; Qin, Y.; Luo, H.; Zhao, F.; Xu, B.; Zeng, L.; and Wang, C. 2023. When Spatio-Temporal Meet Wavelets: Disentangled Traffic Forecasting via Efficient Spectral Graph Attention Networks. In *ICDE*.

Gao, X.; Chen, H.; and Haworth, J. 2023. A spatiotemporal analysis of the impact of lockdown and coronavirus on London's bicycle hire scheme: from response to recovery to a new normal. *Geo-spatial Information Science*.

Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *AAAI*.

Gören, H. G. 2018. A decision framework for sustainable supplier selection and order allocation with lost sales. *J. Clean*.

Hao, S.; Liu, Y.; Wang, Y.; Wang, Y.; and Zhe, W. 2022. Three-stage root cause analysis for logistics time efficiency via explainable machine learning. In *SIGKDD*.

Hasan, M. M.; Jiang, D.; Ullah, A. S.; and Noor-E-Alam, M. 2020. Resilient supplier selection in logistics 4.0 with heterogeneous information. *EXPERT SYST APPL*.

Hu, X.; et al. 2018. Joint decision model of supplier selection and order allocation for the mass customization of logistics services. *Part E*.

Hui, B.; Fang, Y.; Xia, T.; Aykent, S.; and Ku, W.-S. 2023. Constrained market share maximization by signal-guided optimization. In *AAAI*.

Islam, S.; Amin, S. H.; and Wardley, L. J. 2021. Machine learning and optimization models for supplier selection and order allocation planning. *Int J Prod Econ*.

Jiang, X.; Qin, Z.; Xu, J.; and Ao, X. 2023a. Incomplete graph learning via attribute-structure decoupled variational auto-encoder. In *WSDM oral*, 304–312.

Jiang, X.; Qiu, R.; Xu, Y.; Zhang, W.; Zhu, Y.; Zhang, R.; Fang, Y.; Chu, X.; Zhao, J.; and Wang, Y. 2024. RAGraph: A General Retrieval-Augmented Graph Learning Framework. *NeurIPS*.

Jiang, X.; Zhuang, D.; Zhang, X.; Chen, H.; Luo, J.; and Gao, X. 2023b. Uncertainty quantification via spatial-temporal tweedie model for zero-inflated and long-tail travel demand prediction. In *CIKM*, 3983–3987.

Kara, M.; Ulucan, A.; and Atici, K. B. 2019. A hybrid approach for generating investor views in Black–Litterman model. *Expert Syst. Appl.*

Kawtummachai, R.; and Van Hop, N. 2005. Order allocation in a multiple-supplier environment. *Int J Prod Econ*.

Kuroiwa, R.; and Beck, J. C. 2023. Domain-Independent Dynamic Programming: Generic State Space Search for Combinatorial Optimization. In *ICAPS*.

Lea, C.; Flynn, M. D.; Vidal, R.; Reiter, A.; and Hager, G. D. 2017. Temporal convolutional networks for action segmentation and detection. In *CVPR*.

LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *nature*.

Lee, A. H. 2009. A fuzzy supplier selection model with the consideration of benefits, opportunities, costs and risks. *Expert Syst. Appl*.

Li, D.; Liu, J.; Jeon, J.; Hong, S.; Le, T.; Lee, D.; and Park, N. 2021. Large-Scale Data-Driven Airline Market Influence Maximization. In *SIGKDD*.

Li, K.; Liu, Y.; Ao, X.; Chi, J.; Feng, J.; Yang, H.; and He, Q. 2022a. Reliable Representations Make A Stronger Defender: Unsupervised Structure Refinement for Robust GNN. In *SIGKDD*.

Li, R.; Zhong, T.; Jiang, X.; Trajcevski, G.; Wu, J.; and Zhou, F. 2022b. Mining Spatio-Temporal Relations via Self-Paced Graph Contrastive Learning. In *SIGKDD*.

Liu, Y.; Ao, X.; Dong, L.; Zhang, C.; Wang, J.; and He, Q. 2022. Spatiotemporal Activity Modeling via Hierarchical Cross-Modal Embedding. *TKDE*.

Liu, Y.; Ao, X.; Qin, Z.; Chi, J.; Feng, J.; Yang, H.; and He, Q. 2021. Pick and Choose: A GNN-based Imbalanced Learning Approach for Fraud Detection. In *WWW*.

Markowitz, H. 1952. Portfolio Selection. *J Finance*.

Meena, P. L.; and Sarmah, S. P. 2013. Multiple sourcing under supplier failure risk and quantity discount: A genetic algorithm approach. *Part E*.

Mitra, R.; and Zhang, C.-H. 2014. Multivariate Analysis of Nonparametric Estimates of Large Correlation Matrices.

Mossin, J. 1966. EQUILIBRIUM IN A CAPITAL ASSET MARKET. *Econometrica*.

Nasiri, M. M.; Rahbari, A.; Werner, F.; and Karimi, R. 2018. Incorporating supplier selection and order allocation into the vehicle routing and multi-cross-dock scheduling problem. *Int. J. Prod. Res*.

Puvska, A.; and Stojanovic, I. 2022. Fuzzy Multi-Criteria Analyses on Green Supplier Selection in an Agri-Food Company. *JIMD*.

Ren, Y.; Guo, S.; and Sutherland, D. J. 2022. Better supervisory signals by observing learning paths. *arXiv*.

Rezaei, A.; Rahiminezhad Galankashi, M.; Mansoorzadeh, S.; and Mokhatab Rafiei, F. 2020. Supplier selection and order allocation with lean manufacturing criteria: an integrated MCDM and Bi-objective modelling approach. *ENG MANAG J*.

Rosenblatt, F. 1963. PRINCIPLES OF NEURODYNAMICS. PERCEPTRONS AND THE THEORY OF BRAIN MECHANISMS. *AJP*.

Sharpe, W. 1963. A Simplified Model for Portfolio Analysis. *Manage Sci*.

Steinbach, M. C. 2001. Markowitz Revisited: Mean-Variance Models in Financial Portfolio Analysis. *SIAM REV*.

Tan, K.-C.; Kannan, V. R.; and Handfield, R. B. 1998. Supply chain management: supplier performance and firm performance. *Int. J. Purch. Mater. Manage.*

Tibshirani, R. 1996. Regression Shrinkage and Selection Via the Lasso. *J R STAT SOC B*.

Treynor, J. L. 1962. Toward a Theory of Market Value of Risky Assets.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *ICLR*.

Wan, S.-P.; Rao, T.; and Dong, J.-Y. 2023. Time-series based multi-criteria large-scale group decision making with intuitionistic fuzzy information and application to multi-period battery supplier selection. *Expert Syst. Appl*.

Wang, Z.; and Bovik, A. C. 2009. Mean squared error: Love it or leave it? A new look at signal fidelity measures. *IEEE Signal Process Mag*.

Way, R.; Lafond, F.; Lillo, F.; Panchenko, V.; and Farmer, J. D. 2018. Wright meets Markowitz: How standard portfolio theory changes when assets are technologies following experience curves. arXiv:1705.03423.

Wu, Y.; Zhuang, D.; Labbe, A.; and Sun, L. 2021a. Inductive Graph Neural Networks for Spatiotemporal Kriging. In *AAAI*.

Wu, Y.; Zhuang, D.; Lei, M.; Labbe, A.; and Sun, L. 2021b. Spatial Aggregation and Temporal Convolution Networks for Real-time Kriging. *arXiv preprint arXiv:2109.12144*.

Xing, F.; Cambria, E.; Malandri, L.; and Vercellis, C. 2018. Discovering Bayesian Market Views for Intelligent Asset Allocation. *RePEc*.

Yang, K.; Xu, Y.; Zou, P.; Ding, H.; Zhao, J.; Wang, Y.; and Xie, B. 2023. KerPrint: local-global knowledge graph enhanced diagnosis prediction for retrospective and prospective interpretations. In *AAAI*, volume 37, 5357–5365.

Zhang, R.; Jiang*, X.; Fang*, Y.; Luo, J.; Xu, Y.; Zhu, Y.; Chu, X.; Zhao, J.; and Wang, Y. 2024. Infinite-horizon graph filters: Leveraging power series to enhance sparse information aggregation. *arXiv*.

Zhang, T.; Zhang, Y.; Cao, W.; Bian, J.; Yi, X.; Zheng, S.; and Li, J. 2022. Less Is More: Fast Multivariate Time Series Forecasting with Light Sampling-oriented MLP Structures.

Zhao, F.; Xu, Z.; Wang, L.; Zhu, N.; Xu, T.; and Jonrinaldi, J. 2023. A Population-Based Iterated Greedy Algorithm for Distributed Assembly No-Wait Flow-Shop Scheduling Problem. *IEEE Trans Industr Inform*.

Zhong, Q.; Liu, Y.; Ao, X.; Hu, B.; Feng, J.; Tang, J.; and He, Q. 2020. Financial Defaulter Detection on Online Credit Payment via Multi-view Attributed Heterogeneous Information Network. In *WWW*.

Zhuang, D.; Wang, S.; Koutsopoulos, H.; and Zhao, J. 2022. Uncertainty Quantification of Sparse Travel Demand Prediction with Spatial-Temporal Graph Neural Networks. In *SIGKDD*.