

國立臺灣大學電機資訊學院資訊網路與多媒體研究所



碩士論文

Graduate Institute of Networking and Multimedia

College of Electrical Engineering and Computer Science

National Taiwan University

Master's Thesis

新聞特徵融合於多模態深度強化學習
之投資組合最佳化：以台灣股市為例

News Feature Fusion in Multimodal Deep Reinforcement
Learning for Portfolio Optimization: Evidence from the
Taiwan Stock Market

林俊霆

Jun-Ting Lin

指導教授：陳尚澤 博士、張智星 博士

Advisor: Shang-Tse Chen, Ph.D.

Jyh-Shing Roger Jang, Ph.D.

中華民國 115 年 6 月

June 2026

國立臺灣大學碩士學位論文
口試委員會審定書

MASTER'S THESIS ACCEPTANCE CERTIFICATE
NATIONAL TAIWAN UNIVERSITY

新聞特徵融合於多模態深度強化學習之投資組合最佳
化：以台灣股市為例

News Feature Fusion in Multimodal Deep Reinforcement
Learning for Portfolio Optimization: Evidence from the
Taiwan Stock Market

本論文係林俊霆（學號 R13944024）在國立臺灣大學資訊網路
與多媒體研究所完成之碩士學位論文，於民國 115 年 6 月 17 日承下列
考試委員審查通過及口試及格，特此證明。

The undersigned, appointed by the Graduate Institute of Networking and Multimedia on 17 June
2026 have examined a Master's Thesis entitled above presented by JUN-TING LIN (student ID:
R13944024) candidate and hereby certify that it is worthy of acceptance.

口試委員 Oral examination committee:

eSigned by:
陳尚澤
06/18/2026 @ 10:06 GMT+8
(指導教授 Advisor)

eSigned by:
張智星
06/18/2026 @ 17:11 GMT+8

eSigned by:
陳永耀
06/18/2026 @ 16:23 GMT+8

eSigned by:
葉倚任
06/22/2026 @ 14:41 GMT+8

鄭卜壬

系（所）主管 Director:





Acknowledgements

首先，衷心感謝張智星教授，帶領我透過 FinLab 平台實作量化策略，並引薦投資經理人莊顯琛 Alan 與吳康祺 Peter，使我得以在實務視角下踏入量化交易的世界。感謝陳永耀教授，每次會議均耐心聆聽長達半小時至一小時的進度報告，並針對每個環節給予精準的點評，以問答方式引導思考與表達，讓我在口試當天能從容面對口委的提問。感謝陳尚澤教授多次協助審核論文相關文件。

感謝玉山銀行的指導同仁葉倚任博士與林其翰博士，每次會議中提出具體且深入的建議，使本篇碩論在深度與嚴謹性上得以持續提升。

在實驗室的學習歷程中，感謝莊謹譽（Morris）學長在技術與待人處世上的傾囊相授，並協助設置 GPU Server、帶領我共同發表 IEEE 論文，是此研究得以完成的最重要推手。感謝吳克洋在碩一碩二期間一同修課、合作 project，共同走過每一個挑戰。感謝學妹郭奕里陪伴走過玉山成果發表；感謝黃禹（Tommy）學長的演講邀請與 AI 寫作方法的分享；感謝沈昇廷（Ryan）學長帶領我了解配對交易與期權的相關知識；感謝張開學長每次在 Server 出狀況時的即時協助。

最後，感謝父母與家人在求學過程中給予的全力支持，讓我得以無後顧之憂地專注於研究。也感謝所有朋友與實驗室夥伴，陪我解決問題、面對挑戰。

一篇研究的完成絕非單打獨鬥，謹向所有曾給予協助與鼓勵的師長、同儕與家人，獻上最誠摯的感激。





摘要

隨著大型語言模型的崛起，新聞等非結構化資訊成為量化投資中可提取 alpha 的訊號來源。然而，將新聞情緒整合至多股票深度強化學習（Deep Reinforcement Learning，DRL）投資組合最佳化的研究仍屬稀缺。

本研究以 VisionTrader 為基礎架構，於台股元大台灣 50 成分股上探索新聞情緒融入多股票 DRL 投資組合框架的可行性，圍繞三個遞進目的展開：（一）檢視 MSU（Market Score Unit）的貢獻以確立基準架構；（二）比較不同情緒整合配置在不同覆蓋率股票池的效能，並比較四種融合機制以找出最佳整合方式；（三）跨八個時間區段驗證穩健性。

研究結果表明，MSU 在月頻設定下貢獻有限，可安全移除。融合效益高度依賴股票池新聞覆蓋率，高覆蓋率子集可獲顯著提升，混合覆蓋率池則反而衰退；四種融合機制中 Gate 表現最佳。跨期分析顯示，純價量模型在八個區段中有六個優於 0050；新聞融合相對於純價量的優勢具市場環境依賴性——多頭市場帶來顯著增益，但平盤高雜訊市場此優勢消失。

本研究揭示新聞覆蓋率為融合效益的關鍵調節因素，填補新聞情緒與多股票 DRL 投資組合整合的文獻缺口，並為從業者提供實務指引。

關鍵字：深度強化學習、投資組合最佳化、新聞情緒分析、多模態融合、FinBERT、台灣股市





Abstract

With the rise of large language models, unstructured information such as financial news has emerged as a valuable source of alpha signals in quantitative investing. However, research integrating news sentiment into multi-stock deep reinforcement learning (DRL) portfolio optimization remains scarce.

This study takes VisionTrader as the backbone architecture and explores the feasibility of incorporating news sentiment into a multi-stock DRL portfolio framework on Taiwan's Yuanta Taiwan 50 ETF (0050) constituents, organized around three progressive objectives: (1) examining the contribution of the Market Score Unit (MSU) to establish a baseline architecture; (2) comparing different sentiment integration configurations across stock pools with varying news coverage rates, and comparing four fusion mechanisms to identify the optimal integration method; (3) verifying cross-period robustness across eight rolling time windows.

Results show that MSU contributes marginally under monthly settings and can be

safely removed. Fusion effectiveness is highly dependent on news coverage rate: high-coverage subsets achieve significant improvements, while mixed-coverage pools suffer performance degradation; among the four fusion mechanisms, Gate performs best. Cross-period analysis shows that the pure price-volume model outperforms 0050 in six out of eight segments; the advantage of news fusion over the pure price-volume model is market-environment-dependent—yielding significant gains in bull markets but diminishing in sideways, high-noise markets.

This study reveals news coverage rate as a critical moderating factor for fusion effectiveness, fills the literature gap at the intersection of news sentiment and multi-stock DRL portfolio optimization, and provides practical guidelines for practitioners.

Keywords: Deep Reinforcement Learning, Portfolio Optimization, News Sentiment Analysis, Multimodal Fusion, FinBERT, Taiwan Stock Market



Contents


	Page
Verification Letter from the Oral Examination Committee	i
Acknowledgements	iii
摘要	v
Abstract	vii
Contents	ix
List of Figures	xv
List of Tables	xvii
Chapter 1 Introduction	1
1.1 Research Background	1
1.2 Research Motivation	3
1.3 Research Objectives	4
1.4 Thesis Structure	5
Chapter 2 Literature Review	7
2.1 Portfolio Optimization Theory	7
2.1.1 Markowitz Mean-Variance Model	7
2.1.2 Further Developments and Limitations of Traditional Methods	8
2.2 Reinforcement Learning for Portfolio Management	9

2.2.1	Fundamentals of Reinforcement Learning	9
2.2.2	Policy Gradient Methods	10
2.2.3	Representative Research on Deep Reinforcement Learning for Portfolio Management	11
2.3	DeepTrader, VisionTrader, and TrendTrader	11
2.3.1	DeepTrader: Dual-Unit Architecture for RL Portfolio Management	12
2.3.2	VisionTrader: Enhanced Spatiotemporal Modeling with ViT	13
2.3.3	TrendTrader: A Reinforcement Learning Portfolio Framework with News Sentiment	13
2.4	News Sentiment Analysis: Empirical Foundations and Methodological Evolution	15
2.4.1	Empirical Research on the Impact of News Sentiment on Stock Markets	15
2.4.2	Evolution of Sentiment Indicator Construction Methods	15
2.5	Financial Text Sentiment Analysis: Discriminative and Generative Models	17
2.5.1	BERT Series: Discriminative Models for the Financial Domain	17
2.5.2	Applications of Generative LLMs in Financial Text	18
2.6	Multimodal Fusion: Integrating Price-Volume Data and News Sentiment	19
2.6.1	Related Research on Integrating Structured and Unstructured Information	19
2.6.2	Fusion Timing and Feature-Level Methods in Multimodal Fusion	20


Chapter 3 Methodology 23

3.1	System Architecture	23
-----	-------------------------------	----


3.1.1	ASU Architecture Overview	25
3.1.2	Output and Stock Selection	25
3.2	News Sentiment Feature Extraction	26
3.2.1	News Preprocessing Pipeline	26
3.2.2	FinBERT News Sentiment Score (Path A)	27
3.2.3	FinBERT News Sentiment Embedding (Path B)	28
3.2.4	Temporal Aggregation	29
3.3	Feature Fusion Strategy	30
3.3.1	Fusion Architecture Overview	30
3.3.2	Concatenation	31
3.3.3	Addition	31
3.3.4	Gating	32
3.3.5	Cross-Attention	33
3.3.6	Comparison of Fusion Methods	34
3.4	Reinforcement Learning Training Setup	34
3.4.1	Problem Formalization (MDP)	34
3.4.2	Stock Selection Weight Calculation and Capital Allocation	36
3.4.3	Trade Execution Flow	37
3.4.4	Policy Optimization Objective	38
Chapter 4	Experimental Design	39
4.1	Dataset	39
4.1.1	Stock Data	39
4.1.2	News Data	40



4.1.3	News Preprocessing	43
4.2	Experimental Settings	44
4.2.1	Data Splitting	44
4.2.2	Baseline Models	45
4.2.3	Evaluation Metrics	46
4.2.4	Implementation Details	47
Chapter 5	Experimental Results and Discussion	51
5.1	Baseline Performance	51
5.1.1	Experimental Configuration	51
5.1.2	Performance Comparison	52
5.1.3	Cumulative Wealth Curves	53
5.1.4	Summary	54
5.2	MSU Removal Experiment	54
5.2.1	Experimental Configuration	54
5.2.2	Analysis of MSU Output ρ Characteristics	55
5.2.3	Mixed Pool (47 Stocks) Performance Comparison	56
5.2.4	High-Coverage Pool (22 Stocks) Performance Comparison	57
5.2.5	Training Cost Comparison	58
5.2.6	Summary	59
5.3	ASU News Fusion Experiments	59
5.3.1	Experimental Configuration	59
5.3.2	Mixed Pool (47 Stocks)	60
5.3.3	High-Coverage Pool (22 Stocks)	62



5.3.4	Summary	63
5.4	Comparison of Different Fusion Methods	64
5.4.1	Experimental Configuration	64
5.4.2	Performance Comparison	65
5.4.3	Summary	67
5.5	Robustness Analysis	68
5.5.1	Recent Rolling Window Validation	68
5.5.2	Early Rolling Window Validation	71
5.5.3	Validation Period Sensitivity Analysis	73
5.5.4	Summary	75
Chapter 6	Conclusions and Future Research Directions	77
6.1	Research Conclusions	77
6.2	Research Contributions	79
6.3	Research Limitations	82
6.4	Future Research Directions	84
	References	87
	Appendix A — News Coverage Details for 47 Constituent Stocks	91







List of Figures

Figure 3.1	Multimodal ASU Architecture Diagram	24
Figure 3.2	Architecture Diagram of the Four Feature Fusion Methods	30
Figure 4.1	News Coverage Rate of 47 Constituent Stocks (sorted by coverage rate; blue dashed line indicates 70% threshold)	42
Figure 5.1	A/M/47 vs. A/M/22 Cumulative Wealth Curves (2021–2025)	53
Figure 5.2	A/47 vs. A/M/47 Cumulative Wealth Curves (2021–2025)	56
Figure 5.3	A/22 vs. A/M/22 Cumulative Wealth Curves (2021–2025)	58
Figure 5.4	Mixed Pool News Fusion Cumulative Wealth Curves (2021–2025)	61
Figure 5.5	High-Coverage Pool News Fusion Cumulative Wealth Curves (2021– 2025)	63
Figure 5.6	Cumulative Wealth Curves for the Four Fusion Methods (2021–2025)	67
Figure 5.7	Recent Rolling Window Illustration (Segments A–D)	69
Figure 5.8	Recent Rolling Window Cumulative Wealth Curves (Segments A– D, 2×2 Subplots)	70
Figure 5.9	Early Rolling Window Illustration	71
Figure 5.10	Early Window Cumulative Wealth Curves (2×2 Subplots)	73
Figure 5.11	Validation Period Sensitivity Analysis Illustration	74
Figure 5.12	Validation Period Sensitivity Analysis Cumulative Wealth Curves (Left: with 2022; Right: without 2022)	75






List of Tables

Table 2.1	Comparison of Sentiment Analysis Methods	17
Table 2.2	Comparison of Discriminative and Generative Model Positioning .	19
Table 3.1	Architecture Comparison Between This Study and the Original Vi- sionTrader ASU	25
Table 3.2	FinBERT Sentiment Class to Score Mapping	27
Table 3.3	Design Comparison of the Four Fusion Methods	34
Table 4.1	Excluded Constituent Stocks	40
Table 4.2	Distribution of Stocks and News Articles by Coverage Tier	41
Table 4.3	News Sentiment Score Statistics (per Stock-Day)	43
Table 4.4	Embedding Zero-Vector Statistics	44
Table 4.5	Training/Validation/Test Period Settings	45
Table 4.6	Hardware Environment	47
Table 4.7	Model Hyperparameter Settings	47
Table 4.8	Training Hyperparameter Settings	48
Table 4.9	Trading Settings	48
Table 4.10	News Fusion-Related Settings	49
Table 5.1	Section 5.1 Experimental Configuration. Input features format is “ASU input / MSU input.”	52
Table 5.2	A/M/47 vs. A/M/22 Risk-Return Metrics	52
Table 5.3	Section 5.2 Experimental Configuration	55
Table 5.4	MSU Output ρ Statistics (ASU Only is fixed at 0)	55
Table 5.5	A/47 vs. A/M/47 Risk-Return Metrics	56

Table 5.6	A/22 vs. A/M/22 Risk-Return Metrics	57
Table 5.7	ASU Only vs. ASU+MSU Training Time Comparison	58
Table 5.8	Section 5.3 Experimental Configuration	60
Table 5.9	Mixed Pool News Fusion Experiment Risk-Return Metrics	60
Table 5.10	High-Coverage Pool News Fusion Experiment Risk-Return Metrics	62
Table 5.11	Section 5.4 Fusion Method Experimental Configuration	64
Table 5.12	Risk-Return Metrics for the Four Fusion Methods (High-Coverage Pool, 22 Stocks)	65
Table 5.13	Recent Rolling Window Experimental Configuration (Segments A–D)	69
Table 5.14	Recent Rolling Window Test Period Performance (A vs. ASE(G) vs. 0050)	69
Table 5.15	Early Rolling Window Experimental Configuration (Segments E–H)	71
Table 5.16	Early Window Test Period Performance (A vs. 0050)	72
Table 5.17	Validation Period Setting Comparison	73
Table 5.18	Validation Period Sensitivity Analysis Test Period Performance (2024– 2025)	74
Table A.1	News Coverage Details for 47 Constituent Stocks (2013–2025)	92





Chapter 1 Introduction

This chapter explains the background, research motivation, objectives, and thesis structure of this study. Starting from the limitations of structured data in traditional quantitative investment, it identifies opportunities for integrating unstructured information in the era of large language models (LLMs) and establishes the gap this study fills in the existing literature.

1.1 Research Background

Traditional Quantitative Investment Dominated by Structured Data Traditional portfolio theory is centered on structured data. Markowitz [1] proposed Modern Portfolio Theory (MPT), constructing the efficient frontier from historical returns and covariance matrices, while Black & Litterman [2] further incorporated investor views to correct estimation errors. In quantitative investment practice, the features used by models primarily fall into three categories: **technical** (OHLCV price-volume data), **institutional** (institutional net buying, margin trading), and **fundamental** (P/E ratio, operating margin, and other financial indicators). All of the above exist in numerical, structured formats that are easy to process and model.

The LLM Era: A Breakthrough in Processing Unstructured Information The emergence of BERT in 2018 opened a new era of pre-trained language models, enabling machines to effectively understand the semantic context of long texts for the first time; subsequently, the rise of the GPT series further expanded the application boundaries of large language models (LLMs). In the financial domain, **text-based unstructured information** such as news, earnings call transcripts, and social media was previously treated as noise because it was difficult to quantify, but is increasingly regarded in the LLM era as a valuable signal for extracting alpha. Exemplified by BloombergGPT [3], domain-specific LLMs trained on financial text have demonstrated high efficiency in tasks such as sentiment analysis and news classification.

Active Industry Adoption Financial institutions and hedge funds have actively integrated NLP and LLM technology into quantitative trading workflows:

- **Bloomberg (2023):** Released BloombergGPT, a 50-billion-parameter domain-specific LLM trained on 363 billion tokens of financial text, applied to news classification, sentiment analysis, and financial question answering [3].
- **Man Group (2025):** The quantitative arm Man Numeric developed AlphaGPT, which uses LLM agents to automatically mine historical data, write strategy code, and conduct backtests; dozens of AI-generated signals have passed review and entered live trading.¹
- **Widespread Use of Alternative Data:** Academic research shows that a growing number of financial institutions and hedge funds have incorporated alternative data

¹Bloomberg (2025, July 10). *Man Group Says Agentic AI Is Now Devising Quant Trading Signals.* <https://www.bloomberg.com/news/articles/2025-07-10/man-group-says-agentic-ai-is-now-devising-quant-trading-signals>

—including news text, social media, satellite imagery, and other non-traditional data sources—into quantitative models as signal sources beyond traditional price-volume data [4].



The above industry trends indicate that the integration of unstructured information has evolved from experimental exploration to a core component of mainstream quantitative investment practice.

1.2 Research Motivation

Scope Limitations of Existing Research The impact of news sentiment on stock markets has been well-documented in the academic literature. Tetlock [5] was the first to empirically confirm the significant relationship between media sentiment and stock market returns; subsequent research further combined deep learning, using domain-specific language models such as FinBERT [6] to extract sentiment features for predicting individual stock price movements [7] or constructing sentiment-based trading strategies [8]. However, these studies almost exclusively focus on **single-asset stock price prediction (classification tasks)**, treating the problem as a static input-output mapping under a supervised learning framework. This fundamentally differs from the problem setting of **multi-stock dynamic portfolio optimization**: the latter requires simultaneously deciding the relative allocation and long-short positions across multiple stocks at each time step, while accounting for interactions among stocks and the overall market state.

Research Gap Research integrating news sentiment features into **deep reinforcement learning (DRL) portfolio optimization frameworks** remains scarce. TrendTrader [9]

is one of the few works attempting this direction, but remains limited in spatiotemporal modeling capacity due to its LSTM-based architecture. VisionTrader [10] demonstrates superior risk-adjusted returns over DeepTrader in cross-market validation on U.S. stocks (DJIA) and Taiwan stocks (TW50) using a pure Vision Transformer architecture, making it the most competitive base architecture in this domain—yet it has not explored the integration of news sentiment features, a clear gap in the existing literature.

Starting Point of This Research This study takes VisionTrader as the foundation and systematically explores the feasibility and benefits of integrating news sentiment features into a multi-stock DRL portfolio framework. The core question is: in end-to-end reinforcement learning portfolio optimization, **can news sentiment information improve risk-adjusted returns? And how do different fusion strategies affect model performance?**

1.3 Research Objectives

This study is organized around three progressive research objectives:

1. Verify the Necessity of VisionTrader Architecture Components

Before introducing news sentiment, establish a robust baseline architecture. The MSU (Market Score Unit) uses market-level indicators to output a long-short capital allocation ratio ρ , theoretically responsible for risk management, but its actual contribution has not been independently verified. This study conducts an ablation experiment by removing the MSU, comparing the performance difference between ASU Only and ASU+MSU across different stock pool sizes, thereby establishing

the baseline model for subsequent experiments.



2. Verify the Effectiveness and Optimality of News Sentiment Features

Integrate news sentiment into the baseline model in multiple ways and compare with the sentiment-free version, answering the core research question: in end-to-end reinforcement learning portfolio optimization, can news sentiment information yield quantifiable improvements in risk-adjusted returns? Specifically, three configurations are compared: news sentiment score (Path A, early fusion), news sentiment embedding (Path B, intermediate fusion), and both combined; experiments are conducted separately on the full 47-stock pool and the 22-stock high-news-coverage subset to investigate the impact of **news coverage rate** on fusion effectiveness. After confirming effectiveness, further compare four fusion mechanisms for intermediate fusion (Path B): Concatenation, Addition, Gate [11], and Cross-Attention [12], to identify the optimal integration method.

3. Verify Cross-Period Robustness

Using a sliding time window across eight market segments (2009–2025) to verify the cross-period robustness of the model, covering diverse market environments such as bull and sideways markets, with 3 independent training runs to confirm result consistency.

1.4 Thesis Structure

This thesis consists of six chapters, with contents as follows.

Chapter 1: Introduction Explains the research background (Section 1.1), motiva-

tion (Section 1.2), objectives (Section 1.3), and thesis structure (Section 1.4). Starting from the limitations of structured data in traditional quantitative investment, it identifies opportunities for integrating unstructured information in the LLM era and explains the gap this study fills in the existing literature.

Chapter 2: Literature Review Reviews six areas relevant to this study: traditional portfolio theory (Section 2.1), reinforcement learning for portfolio management (Section 2.2), the architectural evolution of DeepTrader / VisionTrader / TrendTrader (Section 2.3), financial text sentiment analysis (Section 2.4), domain-specific financial language models (Section 2.5), and multimodal fusion strategies (Section 2.6).

Chapter 3: Methodology Details the technical design of this study, including system architecture and news integration path design (Section 3.1), news sentiment feature extraction pipeline (Section 3.2), multimodal fusion architecture design (Section 3.3), and reinforcement learning training setup and MDP formalization (Section 3.4).

Chapter 4: Experimental Design Describes the experimental data sources and stock pool selection (Section 4.1), evaluation metrics, and hyperparameter settings (Section 4.2).

Chapter 5: Experimental Results and Discussion Presents baseline performance analysis (Section 5.1), MSU removal verification (Section 5.2), ASU news fusion experiments (Section 5.3), comparison of different fusion methods (Section 5.4), and robustness analysis (Section 5.5).

Chapter 6: Conclusions and Future Research Directions Summarizes the main findings of this study (Section 6.1), explains academic and practical contributions (Section 6.2), research limitations (Section 6.3), and proposes possible directions for future research (Section 6.4).



Chapter 2 Literature Review

This chapter reviews six areas relevant to this study: traditional portfolio theory, reinforcement learning for portfolio management, the evolution of representative DRL portfolio architectures, news sentiment analysis methods, domain-specific financial language models, and multimodal fusion strategies. These provide the theoretical and practical foundations for the methodology design in Chapter 3.

2.1 Portfolio Optimization Theory

This section reviews the theoretical foundations of portfolio optimization, beginning with the mean-variance framework established by Markowitz, tracing subsequent theoretical developments and the limitations of traditional methods, and providing the theoretical background for this study's shift toward deep reinforcement learning approaches.

2.1.1 Markowitz Mean-Variance Model

Markowitz [1] proposed **Modern Portfolio Theory (MPT)**, which for the first time mathematically formalized the risk-return tradeoff of a portfolio. Its core assumption is that investors are rational mean-variance optimizers: given expected return vector μ and covariance matrix Σ for each asset, find the weight vector $\mathbf{w}$ that minimizes portfolio

variance:

$$\min_{\mathbf{w}} \mathbf{w}^\top \Sigma \mathbf{w} \quad \text{s.t.} \quad \mathbf{w}^\top \boldsymbol{\mu} = \mu^*, \quad \mathbf{w}^\top \mathbf{1} = 1 \quad (2.1)$$

Solving for different target returns μ^* yields the **Efficient Frontier**—the set of portfolios with the highest return for a given risk level (or lowest risk for a given return). MPT laid the theoretical foundation for quantitative portfolio management and inspired numerous subsequent extensions.

2.1.2 Further Developments and Limitations of Traditional Methods

Further Developments The main drawback of the Markowitz model is its extreme sensitivity to historical estimates: small estimation errors in $\boldsymbol{\mu}$ can cause the optimization result to shift dramatically. Black & Litterman [2] proposed the **Black-Litterman model**, using market-equilibrium implied returns as a prior and updating expected returns in a Bayesian manner with investor views, significantly improving sensitivity to estimation errors and being widely adopted in institutional investment practice. Another line of development is **Risk Parity** [13], which abandons the goal of maximizing returns and instead allocates based on equal risk contribution of each asset to total portfolio risk, exhibiting greater robustness under certain market conditions.

Common Limitations However, the above traditional methods all face the following fundamental limitations:

- **Static Allocation:** Optimization is solved at a fixed point in time and cannot be dynamically adjusted in real time as markets evolve.
- **Reliance on Historical Estimates:** $\boldsymbol{\mu}$ and Σ are estimated from historical data, with

limited reliability in non-stationary financial markets.

- **Normality Assumption:** Measuring risk with variance implies a symmetric distribution assumption, which fails to capture fat-tail risks and extreme events.
- **Inability to Incorporate Unstructured Information:** Cannot integrate non-numerical market signals such as news and sentiment.



These limitations prompted researchers to turn to reinforcement learning methods capable of dynamically learning decision strategies, as detailed in Section 2.2.

2.2 Reinforcement Learning for Portfolio Management

This section introduces the fundamental concepts and representative methods of reinforcement learning applied to portfolio management, covering the Markov Decision Process framework, Policy Gradient algorithms, and early deep reinforcement learning portfolio research, laying the groundwork for the discussion of subsequent DRL architecture developments.

2.2.1 Fundamentals of Reinforcement Learning

Reinforcement Learning (RL) is a machine learning paradigm that learns decision-making strategies through interaction with an environment. Its core framework is the **Markov Decision Process (MDP)** [14], composed of the following elements:

- **Agent:** The entity that learns and executes decisions (i.e., the investment strategy model)

- **Environment:** The external system in which the agent operates (i.e., the financial market)
- **State s_t :** The observation of the environment at time step t (e.g., stock prices, technical indicators)
- **Action a_t :** The decision taken by the agent in state s_t (e.g., portfolio weights)
- **Reward r_t :** The feedback signal from the environment after taking an action (e.g., return rate)



The agent's goal is to learn a policy $\pi(a | s)$ that maximizes the long-term discounted cumulative reward $\mathbb{E}[\sum_t \gamma^t r_t]$. The MDP assumes that the current state contains all information necessary for decision-making (Markov property), enabling sequential decision problems to be formally solved.

2.2.2 Policy Gradient Methods

Portfolio management involves a **continuous action space** (asset weights are real numbers), making traditional value-function methods (e.g., Q-Learning) difficult to apply directly. Policy Gradient methods [15] directly parameterize the policy π_θ , maximizing expected returns via gradient ascent:

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_t \nabla \log \pi_\theta(a_t | s_t) \cdot G_t \right] \quad (2.2)$$

where G_t is the cumulative return from time step t . **REINFORCE** [15] is the foundational Policy Gradient algorithm, updating gradients using Monte Carlo estimates from complete trajectories, requiring no environment model (model-free), and is suitable for dynamic

systems such as financial markets that are difficult to model precisely.



2.2.3 Representative Research on Deep Reinforcement Learning for Portfolio Management

Jiang et al. [16] proposed the **EIIE (Ensemble of Identical Independent Evaluators)** framework, an early representative study applying deep reinforcement learning to portfolio management. Its core design uses an independent neural network evaluator (CNN, LSTM, etc.) with identical architecture for each asset, outputting scores for each asset, which are then converted to portfolio weights via softmax, with log returns as rewards and Policy Gradient for end-to-end training. EIIE demonstrated the potential of DRL for dynamic portfolio adjustment in cryptocurrency markets.

However, EIIE has two main limitations: **first**, individual asset evaluators are independent and cannot capture spatial dependencies among stocks; **second**, it lacks modeling of the overall market state, making it difficult to adjust long-short exposure when systemic market risk increases. These two gaps are precisely the improvement directions targeted by DeepTrader [17], as detailed in Section 2.3.

2.3 DeepTrader, VisionTrader, and TrendTrader

This section traces three deep reinforcement learning portfolio frameworks most directly relevant to this study: DeepTrader pioneered the dual-unit architecture (ASU + MSU), VisionTrader enhanced spatiotemporal modeling with Vision Transformer, and TrendTrader explored the initial integration of news sentiment, together constituting the architectural evolution context of this research.

2.3.1 DeepTrader: Dual-Unit Architecture for RL Portfolio Management



Wang et al. [17] proposed **DeepTrader**, which formulates portfolio management as an MDP and optimizes investment strategies end-to-end with Policy Gradient, while for the first time integrating market state information into the RL framework through two complementary units:

ASU (Asset Scoring Unit) uses GCN to model spatial dependencies among stocks, combined with temporal attention to capture individual stock dynamics, outputting scores for each stock and training the stock selection strategy with log cumulative returns as rewards; experiments show that constructing the stock relationship graph with a causal structure outperforms correlation and industry classification graphs.

MSU (Market Score Unit) takes macroeconomic market indicators as input, extracts market states with LSTM, and outputs a long-short capital allocation ratio ρ using a Gaussian policy, trained with negative maximum drawdown (negative MDD) as the reward, aiming to actively reduce long exposure to avoid downside risk when markets deteriorate.

The two units serve distinct roles—ASU optimizes returns, MSU minimizes risk—integrated into a single loss function for end-to-end training, forming a risk-return balancing mechanism. DeepTrader’s main limitations are that GCN relies on static graph structures, LSTM has difficulty capturing very long-term dependencies, and it does not incorporate unstructured information.

2.3.2 VisionTrader: Enhanced Spatiotemporal Modeling with ViT




Chuang et al. [10] proposed **VisionTrader**, which retains the DeepTrader dual-unit architecture ($ASU + MSU + \rho$) while comprehensively replacing GCN and LSTM with **Vision Transformer (ViT)**. Its core insight is to conceptualize the spatiotemporal structure of financial time-series data as images evolving over time, leveraging ViT's self-attention mechanism to model dynamic correlations among stocks. Ablation experiments show that removing Spatial Attention and replacing the old architecture with dual ViTs (one each for ASU and MSU) yields the best performance, as the hybrid attention mechanism of GCN and LSTM creates conflicts.

VisionTrader also proposes a four-dimensional robustness evaluation framework—Validation-Test consistency, multi-run stability, temporal segment robustness, and cross-market reliability—and demonstrates significant improvements over DeepTrader in ASR, ARR, and MDD in cross-market validation on U.S. stocks (DJIA) and Taiwan stocks (TW50). **This study takes VisionTrader as its base architecture** to explore the integration of news sentiment features (see Chapter 3).

2.3.3 TrendTrader: A Reinforcement Learning Portfolio Framework with News Sentiment

Ding et al. [9] proposed **TrendTrader**, which follows the DeepTrader line of development, focusing on integrating news sentiment into the RL portfolio framework and addressing the training instability caused by multimodal fusion.



Fusion Design (Early Fusion + Intermediate Fusion) TrendTrader employs both fusion approaches simultaneously. **Early Fusion:** Each stock’s sentiment score sequence is directly concatenated with OHLCV and jointly fed into the main path (Feature Mapping $\rightarrow$ LSTM $\rightarrow$ Temporal Attention), outputting per-stock context c_i . **Intermediate Fusion (Embedding Aggregation):** The sentiment sequence takes an independent path, also passing through Feature Mapping $\rightarrow$ LSTM $\rightarrow$ Temporal Attention, producing a single **global sentiment vector $\mathbf{e}$** (an aggregate representation of all constituent stocks’ sentiment), which is then added to each stock’s context:

$$c_i^h = c_i + \beta \cdot \mathbf{e} \quad (2.3)$$

This design allows market-level sentiment information to be superimposed on individual stock representations as a “background tone,” avoiding the noise sensitivity caused by directly concatenating heterogeneous features. TrendTrader also abandons the independent MSU, using the global sentiment vector as a proxy for market-level information.

Trend Heuristic Introduces a Trend Generator to predict each stock’s upward probability $\hat{y}$, injecting it simultaneously into the Portfolio Generator (concatenated to the input) and the reward function (as a portfolio potential term), guiding capital allocation toward stocks with upward trends. Ablation experiments show that the Trend Heuristic is the most impactful innovation in TrendTrader.

Incremental Training (Two-Stage Training) Supervised learning is used to pre-train the Trend Generator and backbone, followed by end-to-end RL fine-tuning initialized from the pre-trained parameters, improving convergence difficulties in high-dimensional finan-

cial RL environments.



2.4 News Sentiment Analysis: Empirical Foundations and Methodological Evolution

This section reviews the empirical literature on the impact of news sentiment on stock markets, and traces the evolution of sentiment indicator construction methods from lexicon-based approaches to deep learning, providing the methodological basis for this study's use of FinBERT for sentiment feature extraction.


2.4.1 Empirical Research on the Impact of News Sentiment on Stock Markets

Tetlock [5] used the daily proportion of pessimistic words in a Wall Street Journal column as a sentiment indicator, finding that high pessimism could significantly predict next-day stock price declines and increased trading volume, with reversals within a week, suggesting the market exhibits sentiment overreaction. This study established the empirical foundation for “media text sentiment → market reaction,” became an important milestone in the field of financial sentiment analysis, and inspired a large body of subsequent research predicting stock market returns from news text.

2.4.2 Evolution of Sentiment Indicator Construction Methods

Foundation of Supervised Machine Learning Framework (2002) Pang et al. [18] used SVM and Naive Bayes with TF-IDF features for movie review sentiment classifica-

tion, establishing the supervised learning framework of “sentiment analysis= text classification,” which became the foundation for all subsequent sentiment analysis methods.



Early Financial Applications and Limitations of General-Purpose Dictionaries (2007)

Tetlock [5] applied the Harvard General Inquirer general-purpose dictionary to financial text analysis, verifying the predictive power of sentiment signals for markets. However, general dictionaries exhibit systematic misclassifications in financial contexts: financial terms such as “liability” and “tax” are incorrectly labeled as negative words.

Improvement with Finance-Specific Dictionaries (2011) Loughran & McDonald [19]

addressed these misclassifications by re-annotating a corpus of SEC 10-K financial reports, constructing a finance-specific sentiment dictionary covering six word categories including positive, negative, and uncertain, significantly improving the accuracy of financial text sentiment classification.

Extension to Chinese (2015) Chang [20] used the Loughran-McDonald dictionary as a base, cross-translated multiple Chinese-English financial dictionaries and validated against Yahoo Chinese financial news, filling the gap in Chinese financial sentiment resources.

Comparison of Methods Table 2.1 summarizes the main differences among the three categories of methods.

Table 2.1: Comparison of Sentiment Analysis Methods

Method	Representative	Advantages	Disadvantages
Supervised ML	Pang et al. (2002)	Flexible framework	Requires annotated corpus; difficulty capturing context
Financial Dictionary	Loughran & McDonald (2011)	Interpretable; no annotation needed	Limited vocabulary; difficulty handling contextual irony
Deep Learning/LLM	FinBERT, Qwen	Strong contextual understanding	High computational cost (see Section 2.5)

2.5 Financial Text Sentiment Analysis: Discriminative and Generative Models

Deep learning has enabled contextual understanding of financial text, forming two parallel technical approaches: **discriminative models** based on BERT (suited for classification) and **generative LLMs** (suited for summarization and reasoning). This study employs both: FinBERT handles sentiment classification (Path A news sentiment score + Path B news sentiment embedding), and Qwen2.5-14B-Instruct handles daily news summarization.

2.5.1 BERT Series: Discriminative Models for the Financial Domain

The pre-trained language model BERT introduced bidirectional attention mechanisms, greatly improving the ability to capture contextual meaning, replacing traditional bag-of-words features. Huang et al. [6] built on this foundation, conducting domain pre-training with extensive financial report and news corpora, proposing **FinBERT**, specifically designed for financial text sentiment classification (positive/neutral/negative). Compared to general BERT, FinBERT significantly improves classification accuracy in finan-

cial contexts, becoming the standard baseline for subsequent financial NLP research. For Chinese, yiyanghkust/finbert-tone-chinese is a version fine-tuned on Chinese financial news based on the FinBERT architecture, extending sentiment classification capability to Chinese corpora, and is adopted in this study.



2.5.2 Applications of Generative LLMs in Financial Text

Unlike the discriminative architecture of the BERT series, generative large language models (LLMs) generate text autoregressively, with capabilities for summarization, reasoning, and instruction following, suited for handling unstructured long-form news.

Wu et al. [3] trained **BloombergGPT** on Bloomberg's proprietary financial corpus, verifying the advantages of domain-specialized LLMs on financial tasks and establishing the feasibility of using LLMs for financial text analysis. Yang et al. [21] proposed the open-source version **FinGPT**, adopting a data-centric approach combined with Low-Rank Adaptation (LoRA) fine-tuning, lowering the training threshold for domain LLMs and enabling the research community to build financial LLMs without proprietary data. Kirtac & Germano [8] directly used general-purpose LLMs to generate sentiment scores and integrated them into trading strategies, demonstrating the contribution of LLM sentiment signals to excess returns.

Table 2.2 summarizes the positioning comparison between the two technical approaches.

Table 2.2: Comparison of Discriminative and Generative Model Positioning

Approach	Representative Model	Task Type	Applicable Scenarios
Discriminative (BERT)	FinBERT	Sentiment Classification	Fast annotation, stable probability distribution output
Generative (LLM)	BloombergGPT, FinGPT	Summarization, Reasoning	Long document understanding, unstructured news preprocessing

2.6 Multimodal Fusion: Integrating Price-Volume Data and News Sentiment

This section reviews related research on integrating structured price-volume data with unstructured news sentiment, and surveys commonly used feature-level fusion methods, providing the theoretical basis for this study’s design of four fusion mechanisms.

2.6.1 Related Research on Integrating Structured and Unstructured Information

Several representative approaches exist for integrating price-volume time-series data with news sentiment information. Chen et al. [7] encoded daily news headlines with FinBERT and concatenated them at the feature level with stock price time-series features processed by LSTM, forming a deep fusion model for stock price prediction, verifying the effectiveness of combining structured and unstructured features at the feature level. Hong [22] used GPT models to generate news sentiment scores, compressed semantic vectors with PCA and combined them with technical indicators, training machine learning ensemble models for Taiwan stock up/down prediction, demonstrating the potential of generative LLM sentiment features in multi-feature integration.

However, both studies focus on **single-asset stock price prediction (classification)**, and have not explored how to integrate news sentiment features into neural network portfolio optimization frameworks **operating on multiple stocks** (such as VisionTrader [10]). This represents a research gap in the existing literature and is the entry point of this study.


2.6.2 Fusion Timing and Feature-Level Methods in Multimodal Fusion

Classification by Fusion Timing Baltrusaitis et al. [23] classify multimodal fusion into three categories by integration timing:

- **Early Fusion:** Directly concatenates raw features from different modalities at the input layer before feeding into a shared model. Simple to implement, but requires features from all modalities to be in the same semantic space.
- **Intermediate Fusion:** Each modality first passes through an independent encoder, then integrates at the embedding layer; this is the most common architecture in current deep learning multimodal models. Chen et al. [7] use this architecture, encoding stock price features with LSTM and news with FinBERT, generating embeddings from each before fusion.
- **Late Fusion:** Each modality completes its full decision path independently, then is weighted and integrated at the decision layer.

Path A in this study corresponds to early fusion, and Path B corresponds to intermediate fusion, as detailed in Section 3.2.

Representative Algorithms for Intermediate Fusion In intermediate fusion architectures, the design of the feature integration mechanism is the core issue. The main methods are as follows:



Concatenation and Addition are the two most basic strategies. Concatenation preserves complete information from both modalities and typically feeds into a fully connected layer for nonlinear interaction; Addition directly adds the two aligned representations, assuming both modalities are already in the same semantic space, with minimal parameters. Chen et al. [7] use Concatenation to concatenate the LSTM temporal representation with the FinBERT embedding before feeding into a fully connected layer for final prediction.

Gated Multimodal Units (GMU) proposed by Arevalo et al. [11], use learnable gates to dynamically determine the contribution ratio of each modality at each feature dimension. Compared to fixed-weight addition, GMU allows the model to adaptively adjust modality importance based on input, offering greater flexibility and interpretability.

Cross-Attention originates from the Transformer attention mechanism of Vaswani et al. [12]. In cross-modal scenarios, the representation of one modality serves as Query and the other as Key/Value, allowing the model to selectively extract information relevant to the current Query from the target modality. This mechanism can capture fine-grained cross-modal interactions with the strongest expressive power, but also the highest computational complexity.





Chapter 3 Methodology

This chapter details the technical design of this study, describing in sequence the overall system architecture and news integration paths, the news sentiment feature extraction pipeline, the multimodal fusion mechanism design, and the reinforcement learning training setup and MDP formalization.

3.1 System Architecture

This study takes the Asset Score Unit (ASU) of VisionTrader [10] as the foundation and adds two news sentiment feature integration paths, forming a multimodal stock selection architecture. Path A is early fusion: the news sentiment score is directly appended to the raw OHLCV input features, jointly entering the model at the input layer together with price-volume data; Path B is intermediate fusion: news sentiment embeddings obtained by encoding news text with FinBERT are integrated with price-volume features at the ASU's internal feature representation layer through a dedicated fusion module. The two paths can be activated individually or simultaneously, corresponding to the feature extraction design in Section 3.2 and the fusion strategy details in Section 3.3.




Figure 3.1: Multimodal ASU Architecture Diagram

3.1.1 ASU Architecture Overview



Figure 3.1 presents the overall multimodal ASU architecture of this study. Compared to the original VisionTrader, this study adds a news path at both the input end and the feature integration layer, while keeping the ViT backbone structure unchanged. Table 3.1 summarizes the main differences between this study and the original VisionTrader.

Table 3.1: Architecture Comparison Between This Study and the Original VisionTrader ASU

Improvement	Original VisionTrader	This Study
Input Features	OHLCV (5-dim)	OHLCV + news sentiment score (6-dim, Path A)
News Integration	—	FinBERT news sentiment embedding (Path B)
Fusion Method	—	Concat / Add / Gate / Cross-Attn

3.1.2 Output and Stock Selection

The ASU outputs a representation vector $\mathbf{v} \in \mathbb{R}^{128}$ for each stock, which produces scores for each stock after the fusion layer (if Path B is enabled) and the Scoring FC layer.

The Portfolio Generator ranks by score and selects:

- **Long:** the top- G stocks with the highest scores, weight vector ω^+
- **Short:** the bottom- G stocks with the lowest scores, weight vector ω^-

Detailed weight calculation formulas and trade execution procedures are described in Section 3.4.

3.2 News Sentiment Feature Extraction



This section describes how to extract news sentiment scores and news sentiment embeddings from raw news data for use by the ASU module. Raw news is processed sequentially through two stages: LLM summarization and FinBERT analysis, outputting news sentiment scores (Path A, early fusion input) and news sentiment embeddings (Path B, intermediate fusion input), then temporally aggregated to align with the weekly price-volume data.


3.2.1 News Preprocessing Pipeline

Daily news is processed through the following pipeline:

1. **News Collection:** Crawl all news related to the constituent stock for that day.
2. **LLM Summarization:** Use Qwen2.5-14B-Instruct to integrate multiple news articles into a single-sentence summary.
3. **FinBERT Processing:** Perform sentiment analysis (Path A) and vector extraction (Path B) on the summary.

The number of daily news articles varies significantly across stocks and market events; popular stocks may have dozens of reports in a single day. The LLM summarization step serves three purposes: **first**, information consolidation—ensuring each trading day produces one standardized summary, resolving the inconsistency in news quantity; **second**, content filtering—the LLM automatically filters noise during summarization, retaining key information truly relevant to the individual stock; **third**, length unification

—compressing multiple news articles of varying lengths into a fixed format for subsequent FinBERT processing.



3.2.2 FinBERT News Sentiment Score (Path A)

This study uses `yyanghkust/finbert-tone-chinese` for sentiment analysis on Chinese financial text. FinBERT outputs three-class logits, with the predicted class obtained via `argmax`:

$$\text{class} = \arg \max(\mathbf{logits}) \quad (3.1)$$

The mapping between class and sentiment score is shown in Table 3.2.

Table 3.2: FinBERT Sentiment Class to Score Mapping

Class	Sentiment Score (S)
Negative	-1
Neutral	0
Positive	$+1$

The news sentiment score $S \in \{-1, 0, +1\}$ serves as the 6th feature, input to the ASU along with OHLCV:

$$\mathbf{X} = [O, H, L, C, V, S] \quad (3.2)$$

This design is **early fusion**, combining the sentiment signal with price-volume features at the model’s input layer, with both jointly passing through the ASU’s temporal modeling pipeline.

3.2.3 FinBERT News Sentiment Embedding (Path B)



CLS Token Embedding Extract the [CLS] token from FinBERT's last layer as the sentence-level semantic representation:

$$\mathbf{e} = \text{FinBERT}(\text{text})_{[\text{CLS}]} \in \mathbb{R}^{768} \quad (3.3)$$

The CLS token is the special token used by BERT-series models to represent the semantics of an entire text segment; its vector contains global semantic information of the input text.

Projection Layer The raw 768-dimensional embedding contains redundant information and differs in dimensionality from the ASU output (128-dim). It is compressed to 128 dimensions through two trainable projection layers:

$$\mathbf{n} = \text{Proj}(\mathbf{e}) \in \mathbb{R}^{128} \quad (3.4)$$

The specific structure is:

Linear(768 $\rightarrow$ 256) $\rightarrow$ LayerNorm $\rightarrow$ GELU $\rightarrow$ Dropout $\rightarrow$ Linear(256 $\rightarrow$ 128) $\rightarrow$ LayerNorm

This design serves two purposes: **first**, aligning dimensionality with the ASU output $\mathbf{v} \in \mathbb{R}^{128}$, enabling subsequent fusion operations such as Addition and Gating to be performed directly; **second**, learning through end-to-end training to retain features relevant to investment decisions, filtering out irrelevant semantic information. The compressed news sentiment embedding $\mathbf{n}$ and the ASU-output price-volume vector $\mathbf{v}$ are fused at the embedding layer, constituting **intermediate fusion**, as detailed in Section 3.3.

Zero-Vector Padding When a stock has no news on a given trading day, it is padded with a zero vector:

$$\mathbf{e}_t = \mathbf{0} \in \mathbb{R}^{768}, \quad \text{if no news on day } t \quad (3.5)$$



3.2.4 Temporal Aggregation

News data is at daily frequency and needs to be aligned with the ASU's weekly price-volume data. Temporal aggregation proceeds in two stages:

Stage 1: Intra-Week Mean (Data Preprocessing) The embeddings from 5 trading days within a week are averaged to obtain a single weekly embedding:

$$\bar{\mathbf{e}}_w = \frac{1}{5} \sum_{d=1}^5 \mathbf{e}_{w,d} \quad (3.6)$$

Stage 2: Take Last Week (NewsEncoder) Within a 13-week rolling window, take the 13th week (last week) embedding as the current period's representative (corresponding to the hyperparameter setting `news_aggregation = "last"`):

$$\bar{\mathbf{e}} = \bar{\mathbf{e}}_T, \quad T = 13 \quad (3.7)$$

After the above two-stage aggregation, each stock's news features are compressed from shape `[batch, stocks, 13, 768]` to `[batch, stocks, 768]`, then fed into the projection layer to output $\mathbf{n} \in \mathbb{R}^{128}$.



3.3 Feature Fusion Strategy

This section describes how to fuse the ASU-output price-volume feature $\mathbf{v} \in \mathbb{R}^{128}$ with the NewsEncoder-output news sentiment embedding $\mathbf{n} \in \mathbb{R}^{128}$ to produce the final representation $\mathbf{f} \in \mathbb{R}^{128}$, which is then passed to the Scoring FC layer for outputting individual stock scores.

3.3.1 Fusion Architecture Overview





Figure 3.2: Architecture Diagram of the Four Feature Fusion Methods

Figure 3.2 presents the four fusion methods compared in this study: Concatenation,

Addition, Gating, and Cross-Attention. All four methods have the same input dimensions ($\mathbf{v}, \mathbf{n} \in \mathbb{R}^{128}$) and output dimensions ($\mathbf{f} \in \mathbb{R}^{128}$); the differences lie in the mechanism for integrating the two modal features, described in detail in the following subsections.



3.3.2 Concatenation

Concatenation concatenates the two feature types, then projects back to the original dimension through a two-layer fully connected network:

$$\mathbf{f} = \text{MLP}([\mathbf{v} ; \mathbf{n}]) \quad (3.8)$$

The MLP structure is $\text{Linear}(256 \rightarrow 256) \rightarrow \text{LN} \rightarrow \text{GELU} \rightarrow \text{Linear}(256 \rightarrow 128) \rightarrow \text{LN}$.

This method fully preserves information from both feature types and learns their interaction through nonlinear transformation, without assuming that price-volume and news features are already in the same semantic space. Simple in design and highly robust in generalization, this study uses it as the baseline method for intermediate fusion.

3.3.3 Addition

Addition is the lightest fusion method, directly performing element-wise addition of the two feature types followed by normalization:

$$\mathbf{f} = \text{LN}(\mathbf{v} + \mathbf{n}) \quad (3.9)$$

This method has very few parameters (only 2×128 parameters from LayerNorm), but implicitly assumes that the two feature types are already in the same semantic space—if there are large scale or directional differences between them, direct addition may cause information masking.

3.3.4 Gating

Gating uses a learnable gate network to dynamically determine the per-dimension contribution ratio of the two feature types, referencing the Gated Multimodal Units (GMU) proposed by 11 with enhancements. The gate value $\mathbf{g}$ and fusion output $\mathbf{f}$ are computed as follows:

$$\mathbf{g} = \sigma(\text{MLP}([\mathbf{v}; \mathbf{n}])) \in \mathbb{R}^{128} \quad (3.10)$$


$$\mathbf{f} = \text{LN}(\mathbf{g} \odot \mathbf{v} + (1 - \mathbf{g}) \odot \mathbf{n}) \quad (3.11)$$

where $\odot$ is element-wise multiplication; the gate network uses a two-layer MLP structure: $\text{Linear}(256 \rightarrow 128) \rightarrow \text{LN} \rightarrow \text{GELU} \rightarrow \text{Linear}(128 \rightarrow 128) \rightarrow \text{Sigmoid}$.

The motivation for each design element is as follows: **Linear**($256 \rightarrow 128$) reduces the concatenated features' dimensionality, forcing the network to learn a joint compressed representation of the two modalities; **LN** $\rightarrow$ **GELU** stabilizes gradients and provides non-linearity, enabling the gate to learn complex feature interactions rather than linear weighting; **Sigmoid** outputs element-wise gate values $\mathbf{g} \in [0, 1]^{128}$, with each dimension independently controlling the proportion of $\mathbf{v}$ or $\mathbf{n}$ used.

Compared to the original GMU (single-layer linear gate with tanh compression of each modality), this study extends the gate network to a two-layer MLP with LayerNorm

added, while retaining the original features $\mathbf{v}$, $\mathbf{n}$ without compression, to enhance the non-linear expressive power of the gate.



3.3.5 Cross-Attention

Cross-Attention uses the news sentiment embedding as Query, the price-volume vector as Key and Value, and performs cross-stock, cross-modal interaction through multi-head attention, following the architecture of 12:

$$Q = W_Q \mathbf{n}, \quad K = W_K \mathbf{v}, \quad V = W_V \mathbf{v} \quad (3.12)$$


$$\text{Attn} = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V \quad (3.13)$$

$$\mathbf{x} = \text{LN}(\mathbf{n} + \text{Attn}) \quad (3.14)$$

$$\mathbf{f} = \text{LN}(\mathbf{x} + \text{FFN}(\mathbf{x})) \quad (3.15)$$

Using 4 attention heads ($d = 128$), with FFN structure $\text{Linear}(128 \rightarrow 512) \rightarrow \text{GELU} \rightarrow \text{Linear}(512 \rightarrow 128)$, and residual connections.

Here Q, K, V have sequence dimension N (number of stocks), and the attention matrix is $N \times N$, meaning each stock's news sentiment embedding references all N stocks' price-volume features when extracting price-volume information. This design can theoretically capture cross-stock supply chain correlations (e.g., upstream news affecting downstream stock selection); computational complexity is $O(N^2d)$ (not necessarily more expensive than other methods when $N \ll d$), but has the most parameters ($\approx 197\text{K}$), requiring more training data.



3.3.6 Comparison of Fusion Methods

Table 3.3: Design Comparison of the Four Fusion Methods

Method	Core Structure	Fusion Layer Parameters	Characteristics
Concatenation	2-layer MLP	$\approx 65K$	Nonlinear feature interaction; robust generalization
Addition	LayerNorm	≈ 256	Lightest; assumes same semantic space
Gating	2-layer MLP + Sigmoid	$\approx 49K$	Adaptive per-dimension modality weights
Cross-Attn	Multi-head Attn + FFN	$\approx 197K$	Cross-stock, cross-modal attention; complexity $O(N^2d)$; most parameters

The parameter counts in Table 3.3 are for the fusion layer itself, excluding the NewsEncoder projection layer (approximately 230K). $d = 128$ is the hidden layer dimension and N is the number of stocks.

3.4 Reinforcement Learning Training Setup

This section describes the reinforcement learning formalization of the portfolio management problem, including the MDP definition, stock selection weight computation, trade execution flow, and policy optimization objective.

3.4.1 Problem Formalization (MDP)

The portfolio management problem is formulated as a Markov Decision Process (MDP) $M = \langle \mathcal{S}, \mathcal{A}, P, R \rangle$, where the agent makes sequential decisions in the financial market environment to maximize long-term risk-adjusted returns.

State Space $\mathcal{S}$ The state at each time step t is jointly composed of individual stock features and market features:

$$s_t = \{\mathcal{X}_t^s; \mathcal{X}_t^m; \mathcal{E}_t\} \quad (3.16)$$

where each component is defined as:

- $\mathcal{X}_t^s \in \mathbb{R}^{N \times 13 \times F_s}$: price-volume features of N stocks over the past 13 weeks ($F_s = 5$ for pure OHLCV; $F_s = 6$ when Path A is enabled, appending the news sentiment score)
- $\mathcal{X}_t^m \in \mathbb{R}^{13 \times F_m}$: market index features over the past 13 weeks
- $\mathcal{E}_t \in \mathbb{R}^{N \times 768}$: news sentiment embeddings for N stocks (when Path B is enabled; otherwise zero vectors), as detailed in Section 3.2

Action Space $\mathcal{A}$ The action consists of three components:

$$a_t = (\omega_t^+, \omega_t^-, \rho_t) \quad (3.17)$$


where $\omega_t^+ \in \mathbb{R}^N$ is the long position weight (nonzero only for top- G stocks), $\omega_t^- \in \mathbb{R}^N$ is the short position weight (nonzero only for bottom- G stocks), and $\rho_t \in [0, 1]$ is the long-short capital allocation ratio (output by MSU).

Policy Design The overall policy π consists of two independent sub-policies:

ASU Policy π^s : Stock selection policy that computes probabilities for each stock via

softmax, then selects top- G (long) and bottom- G (short) stocks to allocate capital:

$$\pi^s(i | \mathcal{X}^s; \theta) = \frac{\exp(v^i; \theta)}{\sum_{j=1}^N \exp(v^j; \theta)} \quad (3.18)$$



MSU Policy π^m : Capital allocation policy that outputs $\tilde{\rho} \sim \mathcal{N}(\mu_t, \sigma_t^2)$ using a Gaussian distribution. Samples $\tilde{\rho}$ during training and uses the mean μ_t during inference, clamped to $[0, 1]$:

$$\pi^m(\rho | \mathcal{X}^m; \theta) = \mathcal{N}(\mu_t, \sigma_t^2; \theta) \quad (3.19)$$

3.4.2 Stock Selection Weight Calculation and Capital Allocation

Stock Selection Weights (ASU Output) After the ASU outputs individual stock scores $v \in \mathbb{R}^N$, the top- G stocks are selected for long positions (set V_t^+) and the bottom- G stocks for short positions (set V_t^-), with weights allocated via softmax:

$$\omega_{t,i}^+ = \frac{\exp(v_i)}{\sum_{j \in V_t^+} \exp(v_j)}, \quad i \in V_t^+; \quad \omega_{t,i}^+ = 0, \quad i \notin V_t^+ \quad (3.20)$$

$$\omega_{t,i}^- = \frac{\exp(1 - v_i)}{\sum_{j \in V_t^-} \exp(1 - v_j)}, \quad i \in V_t^-; \quad \omega_{t,i}^- = 0, \quad i \notin V_t^- \quad (3.21)$$

This study sets $G = 4$, holding 4 long and 4 short stocks at each trading step.

ρ Capital Allocation Mechanism (MSU Output) The MSU outputs $\rho_t \in [0, 1]$ to control the capital allocation ratio between long and short positions: $\rho_t = 0$ corresponds to market-neutral (long-short balanced), $\rho_t = 1$ corresponds to a pure long strategy.



3.4.3 Trade Execution Flow

Each trading step is divided into three phases (following VisionTrader [10] Fig. 1):

Preparation Obtain current capital C_t , stock prices P_t , and ω_t^+ , ω_t^- , ρ_t output by ASU and MSU.

Execution First execute the short position: borrow and sell shorted stocks using $(1 - \rho_t)$ proportion of capital, with proceeds recycled into the long position:

$$C_t^- = C_t \times (1 - \rho_t), \quad W_t^- = \frac{C_t^- \times \omega_t^-}{P_t} \quad (3.22)$$

$$C_t^+ = C_t + C_t^-, \quad W_t^+ = \frac{C_t^+ \times \omega_t^+}{P_t} \quad (3.23)$$

where W_t^+ , W_t^- are the number of long and short shares held respectively. The proceeds from short selling (C_t^-) are recycled into the long position, causing long capital to exceed original capital C_t , forming a self-financing structure.

Settlement After holding for one trading period (21 trading days), positions are settled at new prices P_{t+1} :

$$M_t^+ = W_t^+ \times P_{t+1}, \quad M_t^- = W_t^- \times P_{t+1} \quad (3.24)$$

$$C_{t+1} = M_t^+ - M_t^-, \quad R_t = \frac{C_{t+1}}{C_t} \quad (3.25)$$

A transaction fee of 0.1% is charged at each rebalancing (fee = 0.001).

3.4.4 Policy Optimization Objective



The overall strategy is optimized end-to-end with Policy Gradient, with the joint objective function:

$$\nabla J(\theta) = \sum_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{|\tau|} \log\left(\frac{P_{t+1}}{P_t}\right) \nabla \pi_\theta^s + \lambda \sum_{t=0}^{|\tau|} R_t \nabla \log(\pi_\theta^m) \right] \quad (3.26)$$

- **ASU gradient** (first term): Uses log return as signal, guiding stock selection policy π^s to increase allocation to stocks that will rise in the future.
- **MSU gradient** (second term): Converts the episode's overall maximum drawdown (MDD) into reward signal R_t , updating capital allocation policy π^m with REINFORCE. R_t is implemented as

$$R_t = -2 \cdot (\text{MDD} - 0.5) \quad (3.27)$$

making the reward symmetric around 0.

- λ (corresponding to hyperparameter $\gamma = 0.2$ in the code) controls the learning rate balance between the two policies.

Parameters are updated via gradient ascent: $\theta \leftarrow \theta + \eta \nabla J(\theta)$. Detailed hyperparameter settings are in Chapter 4, Section 4.2.4.



Chapter 4 Experimental Design

This chapter describes the data foundation and experimental settings, including stock data sources, news data crawling and preprocessing, as well as data splitting, baseline models, evaluation metrics, and hyperparameter settings.

4.1 Dataset

This section describes the two categories of data used in this study: historical Taiwan stock data and financial news data, including their respective sources, coverage, and preprocessing procedures.

4.1.1 Stock Data

Data Source This study uses historical Taiwan stock market data provided by FinLab¹, using backward-adjusted prices and five basic columns (OHLCV) including trading volume. Backward-adjusted prices use the **earliest historical price as the fixed base**; whenever corporate actions such as dividend adjustments, stock splits, or capital reductions occur, prices after the event (in the forward direction) are proportionally adjusted so that past historical data remain unchanged. The entire price series is thus continuous and can

¹<https://finlab.tw/>

be directly used to calculate period returns without being affected by changes in share capital.



Stock Pool Selection The constituent stocks of the Yuanta Taiwan 50 ETF (0050) serve as the basis, with the constituent list taken from the December 2025 announcement version². Since the data period of this study is January 1, 2009 to December 31, 2025, some constituent stocks have incomplete data during this period; therefore, the following 3 stocks are excluded (see Table 4.1):

Table 4.1: Excluded Constituent Stocks

Ticker	Name	Reason for Exclusion	Data Start Date
3711	ASE Technology Holding	Merger and relisting as holding company	2019/03/27
6669	Wiwynn Corp.	New company IPO listing	2018/12/25
6919	Kang Pei	New company IPO listing	2024/10/04

The final experiment uses 47 constituent stocks with complete data for the period (hereafter referred to as the **mixed pool**), covering 4,175 trading days in total.

4.1.2 News Data


Data Source News data comes from Cnyes³, crawled from January 1, 2013 to December 31, 2025 (3,176 trading days in total), covering news related to 47 constituent stocks and the overall market index.

Keyword Filtering Each stock is matched using the company name or commonly used abbreviation. Some stocks require exclusion keywords to avoid misidentification, for example: Formosa Plastics (1301) excludes “Formosa Petrochemical”, Nan Ya Plastics

²Pocket Securities Yuanta Taiwan 50 ETF Holdings Detail: <https://www.pocket.tw/etf/tw/0050/fundholding/>

³<https://www.cnyes.com/>

(1303) excludes “Nanya Technology”, and Evergreen Marine (2603) excludes “EVA Air”, etc. The filtering logic is “(include keyword) AND NOT (exclude keyword)”, and news titles containing “listing table”, “NAV table”, and similar tabular news are excluded to reduce data noise.



Trading Day Alignment Filtered news is aligned to trading days by publication time (00:00–23:59), with the trading day list based on the trading records of 0050.

News Coverage Rate News coverage rate is defined as the proportion of trading days with news coverage for that stock during the data period:

$$\text{Coverage} = \frac{\text{Days with News}}{\text{Total Trading Days}} \times 100\% \quad (4.1)$$

Figure 4.1 shows the coverage rate distribution of 47 constituent stocks, revealing a clear structural disparity: 22 stocks exceed 70% coverage (high-coverage pool), while the remaining 25 fall below 70% (low-coverage pool). Detailed coverage rates for each stock are listed in Appendix 6.4.

Table 4.2 shows the number of stocks and news articles at each coverage tier.

Table 4.2: Distribution of Stocks and News Articles by Coverage Tier

Coverage Range	No. of Stocks	Stock Proportion	Total News Articles	News Proportion
>70%	22	46.8%	303,780	78.9%
50–70%	9	19.1%	42,759	11.1%
30–50%	12	25.5%	34,502	9.0%
<30%	4	8.5%	3,781	1.0%
Total	47	100%	384,822	100%

The 22 stocks with coverage above 70% account for only 46.8% of all stocks, yet contribute **78.9%** (303,780 articles) of the total 384,822 news articles, indicating that market





Figure 4.1: News Coverage Rate of 47 Constituent Stocks (sorted by coverage rate; blue dashed line indicates 70% threshold)

information is highly concentrated among a small number of popular stocks. Furthermore, the 22 high-coverage stocks have 49,698 trading days with news, averaging 6.1 news articles per stock per day; TSMC (2330) averages 16.1 per day and Hon Hai (2317) averages 11.5, reaching dozens on news-intensive trading days.



4.1.3 News Preprocessing

LLM Daily Summarization Since the number of daily news articles varies significantly, directly feeding them into the model makes input length difficult to standardize and prone to including repetitive or irrelevant content. This study uses Qwen2.5-14B-Instruct to integrate all relevant news for each stock on each trading day into a single-sentence Chinese summary, achieving information consolidation, content filtering, and format unification, as detailed in Section 3.2.

News Sentiment Score (Path A) FinBERT is used to perform sentiment analysis on each daily summary, outputting labels: +1 (positive), 0 (neutral), −1 (negative); trading days without news are uniformly set to 0. Table 4.3 summarizes the sentiment distribution by stock-day ($47 \text{ stocks} \times 3,176 \text{ days} = 149,272 \text{ combinations}$).

Table 4.3: News Sentiment Score Statistics (per Stock-Day)

Scenario	Stock-Days	Negative (−1)	Neutral (0)	Positive (+1)
<i>Mixed pool (47 stocks)</i>				
With news	94,853	16.1%	28.4%	55.5%
No news	54,419	0.0%	100.0%	0.0%
Overall	149,272	10.3%	54.5%	35.2%
<i>High-coverage pool (22 stocks, >70%)</i>				
With news	49,698	15.9%	29.0%	55.1%
No news	20,174	0.0%	100.0%	0.0%
Overall	69,872	11.3%	49.5%	39.2%

Positive sentiment dominates when news is present (approximately 55%), reflecting

the tendency of financial news to feature positive reports; the sentiment distribution of the 22 high-coverage stocks closely resembles the overall distribution.



News Sentiment Embedding (Path B) FinBERT is used to extract the CLS token embedding (768 dimensions) from each daily summary; trading days without news are padded with zero vectors. Table 4.4 shows zero vector statistics.

Table 4.4: Embedding Zero-Vector Statistics

Stock Group	Stock-Days	Zero Vectors	Zero Vector Ratio
Mixed pool (47 stocks)	149,272	54,419	36.5%
High-coverage pool (22 stocks, >70%)	69,872	20,174	28.9%
Low-coverage pool (25 stocks, $\leq 70\%$)	79,400	34,245	43.1%

The zero-vector ratio for the low-coverage pool (43.1%) is significantly higher than for the high-coverage pool (28.9%). Zero-vector embeddings in the fusion module may produce hidden representations that are confounded with actual news features, a characteristic that will affect the results of news fusion experiments as detailed in Section 5.3.

4.2 Experimental Settings

This section describes the experimental settings, including the data splitting scheme for training, validation, and test periods, the selection of baseline models, the definition of evaluation metrics, and the implementation details of model training.

4.2.1 Data Splitting

This study uses a fixed time-segment split, ensuring that training, validation, and test data do not overlap in time, as shown in Table 4.5.

Table 4.5: Training/Validation/Test Period Settings

Period	Time Range	Years	Trading Days	Purpose
Training	2016-01 to 2020-12	5 years	1,224	Model parameter learning
Validation	2021-01 to 2023-12	3 years	729	Checkpoint selection
Test	2024-01 to 2025-12	2 years	485	Final performance evaluation

The 5-year training period provides sufficient historical data for the model to learn market features; the 3-year validation period covers various market states including bull, bear, and sideways markets for model selection; the 2-year test period uses entirely unseen data to evaluate the model's generalization ability.


4.2.2 Baseline Models

This study uses the following two models as comparison baselines:

VisionTrader: The original VisionTrader architecture, containing two core modules, ASU (Asset Score Unit) and MSU (Market Score Unit), responsible for stock selection (input: individual stock OHLCV, output: long-short portfolio weights) and market timing (input: market index OHLC, output: ρ value) respectively. This is the direct comparison baseline for this study, reflecting pure price-volume strategy performance without news information.

Buy-and-Hold (0050 ETF): The Yuanta Taiwan 50 ETF (0050) serves as the benchmark for passive investment strategy, representing the overall performance of the Taiwan stock market and used to measure whether active strategies have excess return capability.

4.2.3 Evaluation Metrics



This study uses monthly trading (annualization factor $N_y = 12$). Metric definitions are as follows.

Symbol Definitions Let the return for period t be $r_t = (W_t - W_{t-1})/W_{t-1}$ (where W_t is the portfolio value at the end of period t), arithmetic mean return $\mu = \frac{1}{T} \sum_{t=1}^T r_t$, and return standard deviation $\sigma = \sqrt{\frac{1}{T} \sum_{t=1}^T (r_t - \mu)^2}$.

Return Metrics Annualized Return Rate (ARR) uses arithmetic annualization:

$$\text{ARR} = \mu \times N_y \quad (4.2)$$

Arithmetic annualization is consistent with DeepTrader's original implementation, and shares the same arithmetic framework as the ASR denominator (AVOL), ensuring internal consistency between the two. Another common definition is geometric annualization (CAGR): $\text{ARR}_{geo} = (W_T/W_0)^{N_y/T} - 1$, which is slightly lower than arithmetic annualization due to Jensen's inequality; as long as all models use the same formula, comparison fairness is not affected.


Risk-Adjusted Return Metrics Annualized Sharpe Ratio (ASR) assumes a risk-free rate of $r_f = 0$:

$$\text{ASR} = \frac{\text{ARR}}{\text{AVOL}} = \frac{\mu \times N_y}{\sigma \times \sqrt{N_y}} = \frac{\mu}{\sigma} \times \sqrt{N_y} \quad (4.3)$$

Risk Metrics Annualized Volatility (AVOL) and Maximum Drawdown (MDD):

$$AVOL = \sigma \times \sqrt{N_y} \quad (4.4)$$

$$MDD = \max_t \left(\frac{\max_{s \leq t} W_s - W_t}{\max_{s \leq t} W_s} \right) \quad (4.5)$$



4.2.4 Implementation Details

Hardware Environment All experiments were conducted on the servers listed in Table 4.6.

Table 4.6: Hardware Environment

Host	GPU	CPU	RAM
205 / 206	NVIDIA RTX 5090	Intel Core i7-14700	64 GB
207	NVIDIA RTX 5080 × 2	Intel Core i7-14700K	64 GB

Software environment: CUDA 13.0–13.2, PyTorch 2.10.

Model Hyperparameters Model hyperparameter settings are shown in Table 4.7.

Table 4.7: Model Hyperparameter Settings

Parameter	Value	Description
Hidden Dim	128	Hidden layer dimension
Num Blocks	4	Number of TCN blocks
Kernel Size	2	Convolution kernel size
Dropout	0.5	Dropout rate
Window Length	13	Input time window (weeks)
Trade Length	21	Trading period (trading days)

Hidden dimension 128 is consistent with the ASU output vector and news projection layer dimensions, ensuring that fusion methods for Path B can be directly aligned (see Section 3.3). Window Length of 13 weeks combined with Trade Length of 21 trading days (approximately one month) allows the model to use the past quarter’s price-volume information as the basis for each rebalancing decision.

Training Settings Training hyperparameter settings are shown in Table 4.8.

Table 4.8: Training Hyperparameter Settings

Parameter	Value	Description
Optimizer	Adam	Optimizer
Learning Rate	10^{-6}	Learning rate
Weight Decay	0.001	Weight decay
Batch Size	37	Batch size
Epochs	1,500	Number of training epochs
γ	0.2	MSU loss weight coefficient (Section 3.4.4)
Random Seed	3 runs	Multiple independent runs averaged to improve reliability

Since deep learning models on different GPUs and CUDA versions still cannot guarantee completely consistent numerical results even with fixed random seeds, this study improves result reliability by averaging multiple independent runs.

Trading Settings Trading-related settings are shown in Table 4.9.

Table 4.9: Trading Settings

Parameter	Value	Description
Trade Mode	Monthly	Monthly trading
G	4	4 long / 4 short stocks
Transaction Fee	0.1%	Transaction cost
Allow Short	True	Short selling allowed

Each trading step holds 4 long and 4 short stocks ($G = 4$), with short-selling proceeds recycled into the long position, forming a self-financing structure (see Section 3.4.3). Allowing short selling enables the model to profit simultaneously from rising and falling stocks. Transaction costs are set to 0.1%, following the original settings of DeepTrader and VisionTrader to ensure comparability with prior studies. The statutory fee rate in Taiwan is 0.1425% commission (each for buy and sell) plus 0.3% securities transaction tax on the sell side; however, most brokers offer commission discounts to active traders, and 0.1% can be regarded as a reasonable estimate after discounts.

News Fusion Settings News fusion-related settings are shown in Table 4.10.

Table 4.10: News Fusion-Related Settings

Parameter	Value	Description
News Embedding Dim	768 $\rightarrow$ 256 $\rightarrow$ 128	FinBERT output compressed via two projection layers
Fusion Method	Concat / Add / Gate / Cross-Attn	See Section 3.3
Window Length	13 weeks	Input time window
News Aggregation	Last	5-day intra-week average, then take 13th week
News Dropout	0.5	News feature Dropout

The 768-dimensional FinBERT embedding is compressed to 128 dimensions through two projection layers, aligned with the ASU output dimensions before fusion; News Aggregation is set to Last, i.e., taking the weekly mean embedding of the last week in the 13-week rolling window as the current period's news representative (see Section 3.2.4).





Chapter 5 Experimental Results and Discussion

This chapter sequentially presents the results and discussion of five sets of experiments: baseline performance establishment, MSU removal verification, ASU news fusion experiments, comparison of different fusion methods, and cross-period robustness analysis.

5.1 Baseline Performance

This section uses the complete VisionTrader architecture (ASU + MSU) to establish a baseline, comparing performance differences between two stock pool configurations: mixed pool (A/M/47) and high-coverage pool (A/M/22). The experimental goal is to clarify the impact of stock pool size on model performance and establish a comparison basis for subsequent ablation and news fusion experiments.

5.1.1 Experimental Configuration

Both configurations use the same VisionTrader architecture. The ASU input features are the weekly Open, High, Low, Close, and Volume (OHLCV) for each constituent stock,

totaling 5 feature dimensions; the MSU input features are the weekly Open, High, Low, Close (OHLC) for the market index, totaling 4 feature dimensions. The experimental period is divided into: training period 2016–2020 (5 years), validation period 2021–2023 (3 years), and test period 2024–2025 (2 years). The experimental configuration is shown in Table 5.1.

Table 5.1: Section 5.1 Experimental Configuration. Input features format is “ASU input / MSU input.”

Symbol	Description	Constituent Stocks	ASU Input / MSU Input
A/M/47	Mixed pool	47 stocks	OHLCV (5-dim) / OHLC (4-dim)
A/M/22	High-coverage pool	22 stocks (coverage >70%)	OHLCV (5-dim) / OHLC (4-dim)

5.1.2 Performance Comparison

Table 5.2 presents the risk and return metrics for both configurations during the validation and test periods.

Table 5.2: A/M/47 vs. A/M/22 Risk-Return Metrics

Validation Period (2021–2023)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/M/47	1.98	47.62%	8.25%	24.07%
A/M/22	1.47	33.24%	18.81%	22.64%
0050	0.39	7.97%	31.56%	20.34%
Test Period (2024–2025)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/M/47	2.18	54.76%	13.14%	25.18%
A/M/22	1.71	43.33%	14.75%	25.26%
0050	1.82	39.28%	14.37%	21.61%

A/M/47 outperforms A/M/22 comprehensively in both validation and test periods. The test period ASR reaches 2.18, 27% higher than A/M/22 (1.71) and approximately 20% above the market index 0050 (1.82); the annualized return rate (ARR) exceeds A/M/

22 by 11.4 percentage points, with lower maximum drawdown (MDD), achieving superior risk-adjusted performance.



A/M/47's advantage comes from a more diverse pool composition, with larger relative strength differentials among individual stocks, allowing the long-short strategy to extract richer ranking signals. In contrast, A/M/22 uses a coverage rate $>70\%$ as its selection threshold, which may exclude certain high-volatility stocks and thus narrow the alpha capture space.

5.1.3 Cumulative Wealth Curves

Figure 5.1 shows the cumulative wealth trends for both configurations during the validation and test periods, with each period's starting point reset to 1.0, and the shaded area representing the min/max range across 3 independent training runs.





Figure 5.1: A/M/47 vs. A/M/22 Cumulative Wealth Curves (2021–2025)

In Figure 5.1, the light gray and dark gray backgrounds indicate the validation period (2021–2023) and test period (2024–2025) respectively, with each period reset to 1.0 at the start, and shading showing the min/max range across 3 independent training runs. The

blue curve is A/M/47, the green curve is A/M/22, and the red curve is the 0050 market benchmark. A/M/47 leads A/M/22 clearly in the validation period and maintains that lead through the end of the test period; both exceed 0050 for most of the time. The shaded ranges from 3 training runs are generally narrow, indicating good training stability.



5.1.4 Summary

The results in this section establish the reasonableness of A/M/47 as the baseline for subsequent experiments, outperforming A/M/22 and the market index 0050 on all metrics. However, the above results come from the complete architecture including MSU; does the MSU's market timing mechanism actually contribute to excess performance, or is the ASU's own stock selection capability already sufficient? Section 5.2 will answer this question through ablation experiments removing the MSU.

5.2 MSU Removal Experiment

This section verifies the actual contribution of the market timing mechanism to overall performance through ablation experiments removing the MSU. The experimental design compares the performance difference between ASU Only (fixed $\rho = 0$) and ASU+MSU (dynamic ρ output) under the same stock pool. If the two perform similarly, it indicates that MSU did not provide meaningful timing benefit under this study's settings.

5.2.1 Experimental Configuration

The experimental configuration is shown in Table 5.3.

Table 5.3: Section 5.2 Experimental Configuration

Symbol	Description	Architecture	MSU Input
A/47, A/22	ASU Only	ASU only, fixed $\rho = 0$	—
A/M/47, A/M/22	ASU+MSU	ASU + MSU, dynamic ρ output	OHLC (4-dim)

ASU Only fixes $\rho = 0$, meaning equal-weight allocation between long and short positions at each period with no market timing; ASU+MSU has the MSU dynamically output $\rho \in [0, 1]$ based on market OHLC, controlling the long-short capital ratio. The experimental period is the same as in Section 5.1.

5.2.2 Analysis of MSU Output ρ Characteristics

Before comparing performance metrics, we first analyze the actual ρ distribution output by MSU to clarify whether its timing behavior functions as expected.

Table 5.4: MSU Output ρ Statistics (ASU Only is fixed at 0)

Metric	Validation Period (2021–2023)		Test Period (2024–2025)	
	A/M/47	A/M/22	A/M/47	A/M/22
Mean ρ	0.131	0.089	0.087	0.038
Max ρ	0.660	0.472	0.465	0.310
Non-zero Rate	76%	63%	74%	42%

As shown in Table 5.4, the MSU's average ρ during the validation period is only 0.131 (A/M/47) and 0.089 (A/M/22), further degenerating to 0.087 and 0.038 in the test period, far below the level that would materially affect long-short allocation. Although the Non-zero Rate shows that MSU does not output zero values entirely, the extremely low mean implies that the actual long-short ratio adjustment at each step is negligible, practically equivalent to ASU Only with fixed $\rho = 0$. This indicates that in the Taiwan stock monthly trading setting, the MSU cannot extract meaningful market style signals from market OHLC data.

5.2.3 Mixed Pool (47 Stocks) Performance Comparison



Table 5.5: A/47 vs. A/M/47 Risk-Return Metrics

Validation Period (2021–2023)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/47	1.92	47.16%	9.56%	24.61%
A/M/47	1.98	47.62%	8.25%	24.07%
0050	0.39	7.97%	31.56%	20.34%
Test Period (2024–2025)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/47	2.13	53.92%	14.29%	25.48%
A/M/47	2.18	54.76%	13.14%	25.18%
0050	1.82	39.28%	14.37%	21.61%

From Table 5.5, while A/M/47 is slightly better than A/47 on all metrics, the test period ASR gap is only 0.05 (+2.3%), and the ARR gap is 0.84 percentage points, both far exceeding 0050 with a difference that is not statistically significant. This result is consistent with the ρ analysis: MSU has barely changed the long-short allocation ratio, so the ASU's stock selection capability is the main performance driver.




Figure 5.2: A/47 vs. A/M/47 Cumulative Wealth Curves (2021–2025)

In Figure 5.2, the light and dark gray backgrounds indicate the validation period

(2021–2023) and test period (2024–2025) respectively, each period reset to 1.0 at the start with shading showing the min/max range across 3 independent training runs. The solid blue line is A/M/47, the dashed blue line is A/47, and the red curve is the 0050 market benchmark. The two blue lines almost completely overlap, nearly indistinguishable to the naked eye, with both consistently and significantly above 0050.

5.2.4 High-Coverage Pool (22 Stocks) Performance Comparison

Table 5.6 summarizes the risk-return metrics for both configurations in the high-coverage pool.

Table 5.6: A/22 vs. A/M/22 Risk-Return Metrics

Validation Period (2021–2023)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/22	1.47	34.82%	21.16%	23.81%
A/M/22	1.47	33.24%	18.81%	22.64%
0050	0.39	7.97%	31.56%	20.34%
Test Period (2024–2025)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/22	1.71	43.21%	14.75%	25.26%
A/M/22	1.71	43.33%	14.75%	25.26%
0050	1.82	39.28%	14.37%	21.61%

The high-coverage pool results are more straightforward: test period ASR and AVOL are identical, with ARR differing by only 0.12 percentage points. MSU neither improved nor worsened performance, further confirming that its timing effect in this study’s settings approaches zero.

In Figure 5.3, the background and shading settings are the same as Figure 5.2. The solid green line is A/M/22, the dashed green line is A/22, and the red curve is the 0050




Figure 5.3: A/22 vs. A/M/22 Cumulative Wealth Curves (2021–2025)

market benchmark. The two green lines similarly almost completely overlap, with both consistently above 0050.

5.2.5 Training Cost Comparison

Table 5.7: ASU Only vs. ASU+MSU Training Time Comparison

Stock Pool	Config	ASU Training Time	MSU Training Time	Total Time
Mixed pool (RTX 5090)	A/47	≈2.58 hr	—	≈2.58 hr
	A/M/47	≈2.58 hr	≈ 2.06 hr	≈4.64 hr
High-coverage pool (RTX 5080)	A/22	≈2.33 hr	—	≈2.33 hr
	A/M/22	≈2.34 hr	≈ 1.95 hr	≈4.29 hr

ASU+MSU uses two-stage training: first train ASU, then fix ASU parameters and train only MSU. As shown in Table 5.7, removing MSU saves approximately 2 hours of training time (2.06 hours for the mixed pool, 1.95 hours for the high-coverage pool), a reduction in training cost with practical value given the lack of significant performance difference.

5.2.6 Summary



Combining the three perspectives of ρ analysis, performance comparison, and training cost, this section concludes: The MSU fails to effectively function as a timing mechanism in the Taiwan stock monthly trading setting; its ρ output approaching zero makes it practically equivalent to the fixed $\rho = 0$ of ASU Only, with no significant performance difference during both validation and test periods. Given equal cost and equivalent performance, subsequent experiments (from Section 5.3) consistently use **ASU Only** as the baseline architecture, exploring the benefits of integrating news sentiment features on top of this foundation.

5.3 ASU News Fusion Experiments

This section progressively introduces news sentiment features on top of the ASU Only baseline architecture, investigating the impact of news sentiment scores and news sentiment embeddings on stock selection performance. The experiments simultaneously compare the mixed pool (47 stocks) and the high-coverage pool (22 stocks) to clarify the key role of news coverage rate on fusion effectiveness.

5.3.1 Experimental Configuration

This section tests four feature combinations, all using Concatenation as the fusion method, to control variables and focus on the effect of “which news features are introduced,” as shown in Table 5.8.

Each configuration is run separately on 47 stocks (A/47, AS/47, AE/47, ASE/47)

Table 5.8: Section 5.3 Experimental Configuration

Symbol	Description	Input Features
A	Pure Price-Volume Baseline	OHLCV (5-dim)
AS	+News Sentiment Score	OHLCV + news sentiment score (−1, 0, +1) (6-dim)
AE	+News Sentiment Embedding	OHLCV + news sentiment embedding (Concat Fusion)
ASE	+News Sentiment Score+News Sentiment Embedding	OHLCV + news sentiment score + news sentiment embedding (Concat Fusion)

and 22 stocks (A/22, AS/22, AE/22, ASE/22), with the same experimental period as the previous section.


5.3.2 Mixed Pool (47 Stocks)

Table 5.9: Mixed Pool News Fusion Experiment Risk-Return Metrics

Validation Period (2021–2023)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/47	1.92	47.16%	9.56%	24.61%
AS/47	1.81	49.60%	8.63%	27.57%
AE/47	1.68	41.15%	11.87%	24.54%
ASE/47	1.87	39.39%	8.95%	21.14%
0050	0.39	7.97%	31.56%	20.34%
Test Period (2024–2025)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/47	2.13	53.92%	14.29%	25.48%
AS/47	2.27	54.88%	13.77%	24.28%
AE/47	1.97	50.96%	13.64%	25.82%
ASE/47	1.66	42.86%	20.88%	26.00%
0050	1.82	39.28%	14.37%	21.61%

From Table 5.9, the news fusion effect in the mixed pool shows a clear divergence. The news sentiment score (AS/47) produces a slight positive effect on the baseline, with test period ASR rising from 2.13 to 2.27 (+6.6%); however, adding news sentiment embedding causes AE/47's ASR to drop to 1.97 (−7.5%), and ASE/47, which includes both,

experiences a significant decline, with test period ASR of only 1.66 (−22%) and MDD deteriorating to 20.88%.



The root cause of this phenomenon lies in the structural difference in news coverage rates: in the mixed pool, only the high-coverage pool (22 stocks) has news coverage exceeding 70%; the remaining low-coverage pool (25 stocks) has no news available on many trading days and is padded with zero-vector embeddings. After processing through the projection layer and fusion layer, these zero vectors are not neutral inputs but form noise that is confounded with actual news features, disrupting ASU’s stock selection judgment. Even though the news sentiment score (discrete values $-1/0/+1$) has some signal quality, introducing a large number of low-quality news sentiment embeddings still offsets its positive benefits.




Figure 5.4: Mixed Pool News Fusion Cumulative Wealth Curves (2021–2025)

In Figure 5.4, light and dark gray backgrounds indicate validation period (2021–2023) and test period (2024–2025) respectively, each period reset to 1.0 at the start with shading showing the min/max range across 3 independent training runs. The dark blue curve is A/47, light blue is AS/47, orange is AE/47, dark red is ASE/47, and red is 0050. A/47 and AS/47 lead during the validation period; both remain ahead in the test period,

with ASE/47 showing the most notable decline; all four configurations remain above 0050.



5.3.3 High-Coverage Pool (22 Stocks)

Table 5.10: High-Coverage Pool News Fusion Experiment Risk-Return Metrics

Validation Period (2021–2023)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/22	1.47	34.82%	21.16%	23.81%
AS/22	1.42	34.86%	22.58%	24.61%
AE/22	1.88	46.55%	9.90%	24.82%
ASE/22	1.97	45.55%	9.07%	23.13%
0050	0.39	7.97%	31.56%	20.34%
Test Period (2024–2025)				
Config	ASR↑	ARR↑	MDD↓	AVOL↓
A/22	1.71	43.21%	14.75%	25.26%
AS/22	1.69	42.17%	12.82%	24.98%
AE/22	1.83	48.50%	8.49%	26.71%
ASE/22	2.12	54.38%	6.99%	25.89%
0050	1.82	39.28%	14.37%	21.61%

When the stock pool is restricted to the high-coverage pool (22 stocks, all with coverage >70%), results show a significant reversal. As shown in Table 5.10, ASE/22 achieves ASR 2.12 in the test period, a 24% improvement over baseline A/22 (1.71), with MDD dramatically improving from 14.75% to 6.99%, a 53% reduction. Even when only introducing the embedding (AE/22), test period ASR improves to 1.83 (+7%). Adding only the news sentiment score (AS/22) performs almost the same as A/22, indicating that in the high-coverage environment, the semantically rich news sentiment embedding is the primary source of contribution. Notably, ASE/22 (22 stocks) test period ASR and ARR are already close to the A/47 (47 stocks) baseline, yet MDD dramatically improves (6.99% vs. 14.29%), showing that high-coverage news quality can compensate for the gap in stock pool size, becoming a key performance driver.

The improvement in the high-coverage pool stems from the consistency of embedding quality: all constituent stocks have high-frequency news data, with a zero-vector ratio (28.9%) far lower than the mixed pool (36.5%), mostly concentrated on a small number of trading days that are insufficient to systematically disrupt learning. Concat Fusion effectively integrates news semantic representations with price-volume features, allowing the model to learn when to trust news signals and when to rely more on technical indicators.




Figure 5.5: High-Coverage Pool News Fusion Cumulative Wealth Curves (2021–2025)

In Figure 5.5, background and shading settings are the same as Figure 5.4. The dark green curve is A/22, light green is AS/22, orange is AE/22, dark red is ASE/22, and red is 0050. In the test period, ASE/22 performs best; AE/22 is second, with final value between ASE/22 and A/22; A/22 and AS/22 perform similarly; all four configurations remain consistently above 0050.

5.3.4 Summary

The experiments in this section reveal the core determinant of news fusion effectiveness: **news coverage rate**. In the mixed pool, zero-vector embeddings from low-coverage stocks introduce systematic noise, causing ASE/47 to decline 22% compared to the base-

line; in the high-coverage pool, high-quality embeddings enable ASE/22 to achieve the best performance, with ASR improving 24% and MDD improving 53%. This result indicates that news fusion is not a universally effective enhancement method; its benefits are highly dependent on the quality of news data in the stock pool.

All of the above experiments use Concatenation as the sole fusion method. Section 5.4 will further compare the effects of the four fusion mechanisms— Concatenation, Addition, Gating, and Cross-Attention— on the high-coverage pool (22 stocks, the best combination ASE/22).

5.4 Comparison of Different Fusion Methods

Based on the results from Section 5.3, ASE/22 (ASU + news sentiment score + news sentiment embedding) performs best in the high-coverage pool (22 stocks). This section fixes this feature combination and further compares the effects of four different fusion mechanisms—Concatenation, Addition, Gating, and Cross-Attention— for integrating news sentiment embeddings into the ASU, to determine the most suitable fusion architecture for this study’s settings.


5.4.1 Experimental Configuration

The experimental configuration is shown in Table 5.11.

Table 5.11: Section 5.4 Fusion Method Experimental Configuration

Symbol	Fusion Method	Description
Concat	Direct Concatenation	Concatenate news and ASU output vectors, then project to shared space
Add	Direct Addition	Element-wise addition of news and ASU output vectors, then normalize
Gate	Gating Mechanism	Learnable gates dynamically determine per-dimension contribution ratio
Cross Attn	Cross-Attention	News sentiment embedding as Query, price-volume vector as Key/Value

All four methods use the same ASE feature combination (OHLCV + Sentiment + News Embedding) and the high-coverage pool (22 stocks), with the same experimental period settings as the previous section.



5.4.2 Performance Comparison


Table 5.12: Risk-Return Metrics for the Four Fusion Methods (High-Coverage Pool, 22 Stocks)

Validation Period (2021–2023)				
Fusion Method	ASR↑	ARR↑	MDD↓	AVOL↓
Concat	1.97	45.55%	9.07%	23.13%
Add	1.81	44.28%	11.82%	24.42%
Gate	1.84	42.93%	10.99%	23.28%
Cross Attn	0.47	10.12%	23.73%	21.43%
0050	0.39	7.97%	31.56%	20.34%
Test Period (2024–2025)				
Fusion Method	ASR↑	ARR↑	MDD↓	AVOL↓
Concat	2.12	54.38%	6.99%	25.89%
Add	1.88	54.54%	6.91%	29.16%
Gate	2.15	64.18%	5.64%	29.87%
Cross Attn	1.68	25.80%	12.41%	15.44%
0050	1.82	39.28%	14.37%	21.61%

From Table 5.12, the four fusion methods show different ranking patterns during validation and test periods. During the validation period, Concat leads (ASR 1.97), with Gate and Add close behind, while Cross Attn falls far behind (ASR 0.47), approaching the 0050 level (0.39).

In the test period, a structural ranking shift occurs: Gate emerges to lead, achieving ASR 2.15, ARR 64.18%, MDD 5.64%— the best across all four methods; Concat maintains robust second place (ASR 2.12); Add’s ARR (54.54%) and MDD (6.91%) are close to Concat but with lower ASR; Cross Attn’s test period ASR (1.68) falls below 0050

(1.82), the only configuration below the market index among the four.



The performance differences can be understood from the design logic of the fusion mechanisms. **Concat** learns nonlinear interactions between the two modal features through fully connected layers, retaining complete information without assuming both are in the same semantic space, thus maintaining good generalization robustness despite its simple design. **Add** assumes that price-volume and news sentiment embeddings can be directly aligned and added; this assumption ensures the fewest parameters after dimension alignment, but if there are scale differences between the two modalities, additive fusion may cause information masking. **Gate**'s gating mechanism learns a dynamic weight for each feature dimension, enabling the model to adaptively decide “when to trust news and when to rely on price-volume” based on input, achieving more precise modality balance in the Taiwan stock market during the test period, ultimately achieving the highest ARR and lowest MDD combination. **Cross Attn**'s cross-stock attention matrix (dimension $N \times N$) theoretically capable of modeling cross-stock signals such as supply chain upstream-downstream relationships, but with the most parameters ($\approx 197K$, four times Gate), the current training data is insufficient to support effective learning; with only 5 years of training data in this study's setting, the cross-stock attention cannot converge to effective representations, causing significant underperformance during the validation period and falling below the market index during the test period.

In Figure 5.6, light and dark gray backgrounds indicate validation period (2021–2023) and test period (2024–2025) respectively, each period reset to 1.0 at the start with shading showing the min/max range across 3 independent training runs. Blue curve is Concat, orange is Add, green is Gate, dark red is Cross Attn, and red is 0050. Concat leads slightly during the validation period, Gate is second, and Cross Attn falls far behind




Figure 5.6: Cumulative Wealth Curves for the Four Fusion Methods (2021–2025)

the other three; in the test period, Gate emerges to lead, with Concat close behind; Cross Attn is the only configuration that falls below 0050.

5.4.3 Summary

The comparison results in this section show that fusion mechanism design has a significant impact on news integration effectiveness. Gate fusion achieves the best performance in the test period (ASR 2.15, MDD 5.64%), with its gating mechanism dynamically balancing the contributions of price-volume and news modalities, suitable for actual market environments where news signal quality varies. Concat, though simple in design, performs robustly in both validation and test periods, serving as a reliable alternative when computational resources are limited. Cross Attn cannot leverage its theoretical advantages in this study's setting due to insufficient data to support its largest parameter count ($\approx 197K$, four times Gate).

Section 5.5 will use the best Gate fusion model (ASE/22 + Gate) and the pure price-volume baseline (A/22) for robustness analysis, verifying the consistency of performance

across different market segments through a rolling window approach.



5.5 Robustness Analysis

This section uses a Sliding Window design across eight different historical segments (2009–2025) to verify the cross-period robustness of model performance. The comparison subjects are the pure price-volume model **A** (ASU Only) and the best news fusion configuration **ASE(G)** (ASU + Sentiment + Embedding, Gate Fusion), both using the high-coverage pool (22 stocks). The analysis is divided into three parts: Part 1 (Section 5.5.1) compares **A** and **ASE(G)** simultaneously across the four recent rolling windows (Segments A–D); Part 2 (Section 5.5.2) extends the window back to 2009 (Segments E–H); since news data dates back only to 2013, this part verifies only ASU Only (Configuration **A**); Part 3 (Section 5.5.3) replaces the validation period settings to examine the sensitivity of test period performance to checkpoint selection strategy.

5.5.1 Recent Rolling Window Validation

This subsection covers four recent time windows (Segments A–D), with test periods spanning AI bull markets and sideways markets (including the 2022 crash)—two distinctly different market environments—to evaluate the robustness of **A** and **ASE(G)** and the market environment dependency of news fusion benefits. Each segment's settings are shown in Table 5.13.

Figure 5.7 uses a schematic to show the time splits for the four rolling windows, with light yellow blocks for training, light blue for validation, and light green for testing, with each window progressively shifted forward at the same 5+3+2 year ratio.

Table 5.13: Recent Rolling Window Experimental Configuration (Segments A–D)

Segment	Training Period	Validation Period	Test Period (Market Environment)
A (2016–2025, Primary)	2016–2020 (5 years)	2021–2023 (3 years)	2024–2025 (AI stable bull, 0050 ASR 1.82)
B (2015–2024)	2015–2019 (5 years)	2020–2022 (3 years)	2023–2024 (AI surge bull, 0050 ASR 2.37)
C (2014–2023)	2014–2018 (5 years)	2019–2021 (3 years)	2022–2023 (sideways with 2022 crash, 0050 ASR 0.11)
D (2013–2022)	2013–2017 (5 years)	2018–2020 (3 years)	2021–2022 (sideways with 2022 crash, 0050 ASR 0.05)





Figure 5.7: Recent Rolling Window Illustration (Segments A–D)

Table 5.14: Recent Rolling Window Test Period Performance (A vs. ASE(G) vs. 0050)

Segment (Test Period)	Config	ASR↑	ARR↑	MDD↓	AVOL↓
A (2024–2025, AI Bull)	A	1.71	43.21%	14.75%	25.26%
	ASE(G)	2.15	64.18%	5.64%	29.87%
	0050	1.82	39.28%	14.37%	21.61%
B (2023–2024, AI Surge)	A	1.43	32.27%	10.68%	22.53%
	ASE(G)	1.54	40.87%	13.28%	26.58%
	0050	2.37	36.85%	3.77%	15.52%
C (2022–2023, Sideways)	A	1.62	38.18%	9.04%	23.64%
	ASE(G)	1.33	25.77%	17.42%	19.93%
	0050	0.11	2.16%	29.02%	20.25%
D (2021–2022, Sideways)	A	1.84	33.60%	8.04%	18.33%
	ASE(G)	1.11	19.14%	14.99%	17.07%
	0050	0.05	1.09%	31.56%	22.86%



From Table 5.14, A and ASE(G) outperform 0050 in most segments (Segment B is an exception due to a strong 0050 rally).

The benefits of news fusion show significant market environment dependency. In recent AI-driven bull markets (Segments A and B), ASE(G)'s ARR exceeds A by 20.97% and 8.60% respectively; however, in sideways segments (Segments C and D), A overtakes ASE(G), with news fusion not only failing to bring benefits but also noticeably worsening MDD. The likely reason is that the 2022 systemic risk (rate hikes, geopolitical tensions) caused large gains and losses to offset each other, with individual stock news sentiment overwhelmed by macroeconomic events, making it difficult to form consistent stock selection signals; the pure price-volume model A is better able to identify the relative strength of individual stocks in this high-noise environment.




Figure 5.8: Recent Rolling Window Cumulative Wealth Curves (Segments A–D, 2×2 Subplots)

Figure 5.8 uses 2×2 subplots to show cumulative wealth curves for the four rolling windows, with each subplot's light and dark gray backgrounds indicating the validation and test periods respectively, each period reset to 1.0 at the start, with shading showing the

min/max range across 3 independent training runs; blue is A, green is ASE(G), and red is 0050. In recent bull segments (Segments A, B), test period ASE(G) is noticeably above A; in sideways segments (Segments C, D), A overtakes ASE(G); ASE(G) remains above 0050 in all four subplots; A's final value is slightly below 0050 in Segment B (AI surge bull) but above 0050 in the other three segments. The 3-run min/max shading remains narrow across all subplots, confirming high consistency across the three independent training runs and ruling out conclusions driven by randomness.

5.5.2 Early Rolling Window Validation

Since Cnyes news data dates back only to 2013, this part verifies only the pure price-volume model A, extending the sliding window back to 2009 to cover diverse market environments including the COVID-19 V-shaped recovery and the 2018 bear market. Each segment's settings are shown in Table 5.15.

Table 5.15: Early Rolling Window Experimental Configuration (Segments E–H)

Segment	Training Period	Validation Period	Test Period (Market Environment)
E (2012–2021)	2012–2016 (5 years)	2017–2019 (3 years)	2020–2021 (post-COVID V-shaped recovery, 0050 ASR 1.35)
F (2011–2020)	2011–2015 (5 years)	2016–2018 (3 years)	2019–2020 (strong bull, 0050 ASR 1.50)
G (2010–2019)	2010–2014 (5 years)	2015–2017 (3 years)	2018–2019 (bear market then recovery, 0050 ASR 1.00)
H (2009–2018)	2009–2013 (5 years)	2014–2016 (3 years)	2017–2018 (bull then bear, 0050 ASR 0.65)

5.5.2 Earlier Window Validation (Segments E–H, ASU Only)




Figure 5.9: Early Rolling Window Illustration

Figure 5.9 uses a schematic to show the time splits for the four early windows, with the same color scheme as Figure 5.7.




Table 5.16: Early Window Test Period Performance (A vs. 0050)

Segment (Test Period)	Config	ASR↑	ARR↑	MDD↓	AVOL↓
E (2020–2021, Post-COVID Recovery)	A	2.54	58.14%	11.06%	22.93%
	0050	1.35	25.39%	15.90%	18.75%
F (2019–2020, Strong Bull)	A	1.62	24.21%	6.26%	14.99%
	0050	1.50	29.91%	17.74%	19.93%
G (2018–2019, Bear then Recovery)	A	1.26	22.43%	13.60%	17.88%
	0050	1.00	13.63%	14.61%	13.65%
H (2017–2018, Bull then Bear)	A	0.82	17.77%	22.47%	21.45%
	0050	0.65	6.69%	11.80%	10.32%

From Table 5.16, A’s ASR outperforms 0050 in all four early test segments, extending the pure price-volume model’s alpha capture capability back to training data from 2009. Notably, the test period 2017–2018’s MDD reaches 22.47%, the highest among all eight segments, and significantly higher than 0050 in the same period (11.80%), indicating room for improvement in the model’s downside risk control in markets with mixed long and short forces.

Figure 5.10 uses 2×2 subplots to show cumulative wealth curves for the four early windows, with background and shading settings the same as Figure 5.8; blue is A (ASU Only) and red is 0050, with no ASE(G) in this part. A’s final value exceeds 0050 in three of the four early segments (E, G, H); in Segment F (2011–2020, test 2019–2020), 0050’s final value is slightly above A; Segment E (2020–2021, upper left) shows the largest out-performance margin; Segment H (2017–2018, lower right) has the widest shading range, reflecting the highest training instability. Overall, the 3-run shading remains narrow across all subplots (H is slightly wider), confirming consistent results across the three indepen-




Figure 5.10: Early Window Cumulative Wealth Curves (2×2 Subplots)

dent training runs, with conclusions unaffected by random seeds.

5.5.3 Validation Period Sensitivity Analysis

The main experiment's validation period (2021–2023) includes the 2022 crash year, which may bias checkpoint selection toward defensive strategies. This part replaces the validation period with 2019–2021, which excludes the 2022 bear market, while shifting the training period accordingly and keeping the test period at 2024–2025, to evaluate the sensitivity of test period performance to checkpoint selection strategy. The validation period setting comparison is shown in Table 5.17.

Table 5.17: Validation Period Setting Comparison

Experiment Config	Training Period	Validation Period	Val. 0050 ASR	Test Period
Main Experiment	2016–2020 (5 years)	2021–2023 (3 years)	0.39 (with bear)	2024–2025
This Part	2014–2018 (5 years)	2019–2021 (3 years)	1.47 (normal bullish)	2024–2025

Figure 5.11 shows the time split comparison between the two validation period settings, both with the same test period (2024–2025) but different validation periods.

5.5.3 Validation Period Sensitivity Analysis (2014–2025)

Test period fixed at 2024–2025; comparing val period with vs. without the 2022 bear market




Figure 5.11: Validation Period Sensitivity Analysis Illustration

Table 5.18: Validation Period Sensitivity Analysis Test Period Performance (2024–2025)

Validation Period Setting	Config	ASR↑	ARR↑	MDD↓	AVOL↓
Main Experiment (val with 2022)	A	1.71	43.21%	14.75%	25.26%
	ASE(G)	2.15	64.18%	5.64%	29.87%
	0050	1.82	39.28%	14.37%	21.61%
This Part (val without 2022)	A	1.48	47.02%	10.42%	31.87%
	ASE(G)	1.71	53.68%	10.00%	31.45%
	0050	1.82	39.28%	14.37%	21.61%

From Table 5.18, ASE(G) outperforms A under both validation period settings, but the lead narrows noticeably in this part: the ARR gap narrows from 20.97% to 6.66%, and the ASR gap decreases from 0.44 to 0.23. Furthermore, switching to a validation period without the 2022 bear market causes A’s AVOL to rise from 25.26% to 31.87%, and ASE(G) from 29.87% to 31.45%, with both showing reduced volatility control capability.

The above comparison points to the same conclusion: **including the 2022 bear year in the validation period has a positive effect on checkpoint selection.** The original validation period’s (2021–2023) mixed bull-bear environment forces the model to balance risk and return, tending to select checkpoints with better volatility control, resulting in lower test period AVOL and more prominent ASE(G) relative advantage. Once a purely bullish validation period (2019–2021) is used, checkpoints lack bear market calibration, causing test period volatility to rise and all relative advantages to shrink accordingly.

Figure 5.12: the left subplot shows the main experiment (val 2021–2023), the right




Figure 5.12: Validation Period Sensitivity Analysis Cumulative Wealth Curves (Left: with 2022; Right: without 2022)

subplot shows this part (val 2019–2021); both test periods are 2024–2025, each period starting independently from 1.0; shading shows the min/max range across 3 independent training runs; blue is A, green is ASE(G), and red is 0050. Both subplots show $ASE(G) > A > 0050$ during the test period, but the gap between ASE(G) and A's final value in the right subplot is noticeably smaller (approximately 0.29 vs. 0.97 in the left subplot), with the lead substantially narrowed. Both configurations maintain narrow 3-run shading bands, showing that the change in validation period design does not affect the reproducibility of results, with all three training runs converging to consistent conclusions.

5.5.4 Summary

1. A vs 0050: Pure Price-Volume Model is Cross-Period Robust (6/8 Wins) A's ASR outperforms 0050 in six out of eight rolling windows, maintaining $ASR > 1.0$ even in sideways markets, supporting that the relative strength signals captured by the pure price-volume ASU have long-term universality.

2. ASE(G) vs 0050: News Fusion Model Also Wins in Most Segments (3/4 Wins) ASE(G) outperforms 0050 in three out of four recent segments (Segment B is an exception

due to a strong 0050 rally), indicating that news fusion still provides positive contributions in most market environments.



3. ASE(G) vs A: Fusion Benefits Are Market-Environment-Dependent In recent bull markets, ASE(G)'s ARR gains over A reach +20.97% (Segment A) and +8.60% (Segment B); in sideways segments (Segments C and D), A overtakes ASE(G), as individual stock news cannot form consistent stock selection signals in high-noise environments. News fusion benefits vary significantly with market conditions and should not be treated as unconditional performance enhancements.



Chapter 6 Conclusions and Future Research Directions

This chapter summarizes the research findings, presenting in order the core experimental findings, academic and practical contributions, research limitations, and directions for future research.

6.1 Research Conclusions

This study takes VisionTrader as the foundational architecture, systematically exploring the feasibility of integrating news sentiment features into a deep reinforcement learning portfolio framework, and verifying through experiments on Taiwan Stock Exchange TW50 constituent stocks. Regarding the core question raised in Section 1.2—whether news sentiment information can improve risk-adjusted returns and how different fusion strategies affect model performance—this study’s answer is: under conditions of sufficient news coverage, integrating news sentiment embeddings into the ASU with Gate fusion can significantly improve risk-adjusted returns in bull market environments; however, this benefit is market-environment-dependent and does not hold under all market conditions. The following summarizes the core conclusions according to the three re-

search objectives (Section 1.3).



Conclusion 1: MSU Contributes Little in Monthly Taiwan Stock Settings and Can

Be Safely Removed The MSU (Market Score Unit) is designed to dynamically output a long-short capital allocation ratio ρ using market OHLC data, but ablation experiments in Section 5.2 show that the learned ρ values generally approach zero (average 0.087 for the mixed pool, 0.038 for the high-coverage pool), essentially equivalent to a fixed $\rho = 0$ market-neutral allocation. After removing MSU, the ASR difference in the mixed pool is only 2.3%, and the high-coverage pool shows identical results, with training time savings of approximately 2 hours. Given equivalent performance, removing MSU reduces architectural complexity and training costs, establishing a more streamlined baseline model for subsequent experiments.

Conclusion 2: News Sentiment Benefits Are Highly Dependent on Coverage Rate;

Gate Fusion Is the Best Strategy Experimental results in Section 5.3 show that the effect of integrating news sentiment into the ASU differs dramatically based on the stock pool's news coverage rate. In the mixed pool (47 stocks), adding news sentiment embeddings (AE/47) causes ASR to decrease by 7.5%; combining news sentiment score and embedding (ASE/47) leads to a large decline (ASR -22% , MDD deteriorating to 20.88%). The root cause is that zero-vector news sentiment embeddings from the low-coverage pool (25 stocks) become noise in the model, offsetting the positive effects of sentiment signals. In contrast, on the high-coverage pool (22 stocks, coverage $> 70\%$), ASE/22 achieves the best performance (Test ASR 2.12, ARR 54.38%, MDD 6.99%), a 24% improvement in ASR and 53% improvement in MDD compared to the pure price-volume baseline. Adding only the news sentiment score (AS/22) has limited effect, indicating that semantically rich

news sentiment embeddings are the primary source of improvement.

Section 5.4 further compares four intermediate fusion mechanisms on the high-coverage pool (22 stocks), with Gate fusion performing best during the test period (ASR 2.15, ARR 64.18%, MDD 5.64%), Concat being a robust second choice; Cross-Attention performs worst due to its largest parameter count ($\approx 197K$) that cannot be effectively learned with the current data volume. **Core conclusion:** News sentiment fusion effectively improves risk-adjusted returns in the high-coverage pool; Gate fusion is the best integration method; the mixed pool should maintain a pure price-volume model.

Conclusion 3: The Model Has Cross-Period Robustness; News Fusion Benefits Are Market-Environment-Dependent Robustness analysis across eight rolling time windows (2009–2025) in Section 5.5 shows that the pure price-volume model A's ASR outperforms 0050 in six out of eight segments (6/8), and ASE(G) outperforms 0050 in three out of four recent segments (3/4). The relative advantage of ASE(G) vs A is market-environment-dependent: in recent AI-driven bull markets (2024–2025, 2023–2024), ASE(G)'s ARR exceeds the pure price-volume baseline by +20.97% and +8.60% respectively; but in sideways high-noise markets (2022–2023, 2021–2022), the pure price-volume model outperforms ASE(G).

6.2 Research Contributions

This section elaborates on the contributions of this study from both academic and practical perspectives, focusing on which literature gaps are filled, what empirical benchmarks are provided for future research, and design guidance for quantitative investment

practice.



Academic Contributions

First Systematic Exploration of News Sentiment Integration in Multi-Stock DRL

Portfolio Frameworks In the existing literature, research on integrating news sentiment into financial models is almost exclusively focused on single-asset stock price classification tasks. This study expands the problem setting to multi-stock dynamic portfolio optimization, filling the literature gap in the “news sentiment $\times$ DRL portfolio” intersection, and providing complete experimental benchmarks for subsequent research, covering architecture design, data processing pipelines, and evaluation protocols.

Revealing the Decisive Impact of News Coverage Rate on Multimodal Fusion Benefits

This study proposes and validates “stock pool news coverage rate” as a key moderating variable for fusion effectiveness, pointing out that zero-vector news sentiment embeddings are the root cause of fusion failure in mixed-coverage scenarios. This finding is unprecedented in the existing multimodal financial learning literature and provides direct guidance for data selection and experimental design in future research.

Providing Reproducible Empirical Benchmarks for MSU Component Necessity and

Fusion Mechanism Selection The ablation experiments in this study independently verify for the first time the actual contribution of MSU in monthly Taiwan stock settings, filling the gap in the DeepTrader/VisionTrader architecture literature. Furthermore, the systematic comparison of four fusion mechanisms—Concatenation, Addition, Gate, and Cross-Attention—under unified experimental conditions clarifies the applicable condi-

tions of each mechanism in financial time-series and reinforcement learning environments, serving as an empirical reference for future research in selecting fusion mechanisms.



Defining Market Environment Conditions for Sentiment Fusion Benefits Rolling validation across eight time segments reveals “under what market conditions news sentiment is effective,” providing a robustness evaluation protocol that can be repeatedly applied to other fusion strategies, enabling future research to design experiments under more precise conditional assumptions.

Practical Contributions

Actionable Design Guidelines for News Sentiment Fusion The experimental conclusions of this study can be directly translated into practical design recommendations, helping quantitative investment practitioners make evidence-based architectural choices when incorporating news sentiment:

1. **Stock Pool Selection:** News coverage rate should be used as a stock pool selection criterion, with a recommended threshold of $> 70\%$; low-coverage stocks should not be included in news fusion models
2. **Fusion Mechanism Selection:** Gate fusion is the first choice; Concat is a robust alternative when computational resources are limited; Cross-Attention should be avoided in small dataset scenarios
3. **Architecture Simplification:** MSU can be safely removed in monthly rebalancing settings, reducing training costs without sacrificing performance

4. **Deployment Timing Awareness:** News sentiment fusion strategies have significant benefits in bull market environments; market state monitoring should be combined to avoid over-reliance on sentiment signals during periods dominated by systemic risk



6.3 Research Limitations

This study has certain areas for improvement in terms of data, model design, and experimental settings, as described below.

Data Limitations

Single Market Validation This study only conducted experiments on Taiwan Stock Exchange TW50 constituent stocks without cross-market validation (e.g., U.S. DJIA or other emerging markets). Taiwan's stock market has specific structural characteristics (electronics-concentrated, high retail investor proportion, Chinese-language news), and whether the experimental conclusions can be generalized to other markets requires further verification.

Uneven News Coverage Rate and Data Time Range Coverage rates vary widely among constituent stocks in the Cnyes news data (12%–100%), causing the low-coverage pool (25 stocks, $\leq 70\%$) to be unable to effectively use news sentiment features, so news fusion experiments are concentrated on the high-coverage pool (22 stocks). Furthermore, Cnyes data dates back only to 2013, making it impossible to verify news fusion models in earlier time windows.

Cross-Attention Limited by Training Data Volume Cross-Attention fusion performs worst in this study, primarily because it has the largest parameter count ($\approx 197K$), and the current training set (5 years of monthly data, approximately 60 steps) is insufficient for the model to effectively learn meaningful cross-stock attention weights. This is a data volume limitation rather than a fundamental flaw of the mechanism itself, and it still has potential under conditions with more sufficient training data.

Model Limitations

MSU Timing Mechanism Not Redesigned This study found that ρ approaches zero when MSU uses market OHLC as input, and chose to remove MSU to simplify the architecture. However, market timing itself remains an important capability in portfolio management; this study only confirmed that the existing MSU design is ineffective without exploring more effective timing mechanisms, leaving an unfilled architectural gap.


Experimental Limitations

Monthly Trading Settings This study follows VisionTrader's monthly (21 trading days) rebalancing setting without exploring weekly or daily frequency effects. Monthly settings limit the model's ability to respond to short-term news events (such as single-day earnings releases), and the immediacy advantage of news sentiment is not fully leveraged.

Generalizability of Conclusions

Market Environment Dependency of News Fusion Benefits Gate fusion benefits are significant in recent AI-driven bull markets (2023–2025) but worse than pure price-volume

models in sideways high-noise markets (2021–2022). The advantages and disadvantages of news sentiment fusion strategies are highly related to market conditions; the feature-level Gate mechanism cannot solve the regime-switching problem at the strategy level, and may require higher-level adaptive mechanisms under different market conditions.



6.4 Future Research Directions

Cross-Market Validation Corresponding to the “single market validation” limitation, future research can extend the same framework to U.S. stocks (DJIA, S&P 500) or other Asian markets, evaluating the cross-market generalizability of news sentiment fusion strategies and clarifying the extent to which Taiwan-specific market structure affects experimental conclusions.

Expanding News Data Sources Corresponding to the “uneven news coverage rate and data time range” limitation, future research can incorporate diverse news sources (such as financial databases, social media, company announcements) to improve news density for low-coverage stocks, enabling larger stock pools to benefit from news fusion, and extending early data coverage to support robustness validation over longer time windows.

Rebuilding MSU Timing Mechanism Corresponding to the “MSU timing mechanism not redesigned” limitation, future research can explore more effective market state input designs to enable MSU to truly learn meaningful timing signals and perform risk management functions across different market environments, filling the architectural gap left by removing MSU in this study.

Finer Trading Frequency Corresponding to the “monthly trading settings” limitation, future research can explore weekly or even daily frequency settings, shortening the re-balancing cycle to more immediately reflect news events (such as earnings releases and major announcements), fully leveraging the short-term alpha signal advantages of news sentiment.

Market Environment Level Strategy Switching Corresponding to the “market environment dependency of news fusion benefits” limitation, the Gate mechanism in this study dynamically adjusts news sentiment embedding weights at the feature level (micro-level), but cannot solve the regime-switching problem at the strategy level. Future work can explore introducing market state detection at a higher level (such as volatility regimes, trend strength indicators), dynamically deciding whether to activate the news fusion module based on macroeconomic market conditions, enabling the model to maintain optimal performance in both bull and sideways markets.






References

- [1] Harry Markowitz. Portfolio selection. The Journal of Finance, 7(1):77–91, 1952.
- [2] Fischer Black and Robert Litterman. Global portfolio optimization. Financial Analysts Journal, 48(5):28–43, 1992.
- [3] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. BloombergGPT: A large language model for finance, 2023.
- [4] Yunchuan Sun, Lu Liu, Ying Xu, Xiaoping Zeng, Yufeng Shi, Haifeng Hu, Jie Jiang, and Ajith Abraham. Alternative data in finance and business: emerging applications and theory analysis (review). Financial Innovation, 10(1):127, 2024.
- [5] Paul C. Tetlock. Giving content to investor sentiment: The role of media in the stock market. The Journal of Finance, 62(3):1139–1168, 2007.
- [6] Allen H. Huang, Hui Wang, and Yi Yang. FinBERT: A large language model for extracting information from financial text. Contemporary Accounting Research, 40(2):806–841, 2023.
- [7] Pinyu Chen, Zois Boukouvalas, and Roberto Corizzo. A deep fusion model for stock

market prediction with news headlines and time series data. Neural Computing and Applications, 36(34):21229–21271, 2024.



[8] Kemal Kirtac and Guido Germano. Sentiment trading with large language models. Finance Research Letters, 62:105227, 2024.

[9] Wei Ding, Zhennan Chen, Hanpeng Jiang, Yuanguo Lin, and Fan Lin. Trend-heuristic reinforcement learning framework for news-oriented stock portfolio management. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 5120–5124, 2024.

[10] Chun-Yu Marvin Chuang, Jun-Ting Lin, Yung-Yaw Chen, Yi-Ren Yeh, Chun-Yuan Huang, and Jyh-Shing Roger Jang. VisionTrader: Portfolio optimization with deep reinforcement learning and Vision Transformer. Manuscript submitted for publication. Department of Computer Science and Information Engineering, National Taiwan University, 2026.


[11] John Arevalo, Tamar Solorio, Manuel Montes-y-Gómez, and Fabio A. González. Gated multimodal units for information fusion, 2017.

[12] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Advances in Neural Information Processing Systems, volume 30, pages 5998–6008, 2017.

[13] Sébastien Maillard, Thierry Roncalli, and Jérôme Teïletche. The properties of equally weighted risk contribution portfolios. The Journal of Portfolio Management, 36(4):60–70, 2010.

- 
- [14] Richard S. Sutton and Andrew G. Barto. Reinforcement Learning: An Introduction. MIT Press, 2nd edition, 2018.
- [15] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine Learning, 8(3–4):229–256, 1992.
- [16] Zhengyao Jiang, Dixing Xu, and Jinjun Liang. A deep reinforcement learning framework for the financial portfolio management problem, 2017.
- [17] Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. DeepTrader: A deep reinforcement learning approach for risk-return balanced portfolio management with market conditions embedding. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 35, pages 643–650, 2021.
- [18] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up? Sentiment classification using machine learning techniques. In Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 79–86, 2002.
- [19] Tim Loughran and Bill McDonald. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. The Journal of Finance, 66(1):35–65, 2011.
- [20] 張津挺. 中文財務情緒字典建構與其在財務新聞分析之應用. Master's thesis, 臺北市立大學, 2015.
- [21] Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. FinGPT: Open-source financial large language models, 2023.
- [22] 洪嘉濃. 整合大型語言模型情感分析與機器學習之股票漲跌預測方法. Master's thesis, 國立臺灣大學, 2025.

- [23] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. IEEE Transactions on Pattern Analysis and Machine Intelligence, 41(2):423–443, 2019.





Appendix A — News Coverage Details for 47 Constituent Stocks

Table A.1 lists the news coverage details for the 47 experimental constituent stocks during the data period (January 1, 2013 to December 31, 2025, totaling 3,176 trading days), sorted in descending order by coverage rate.

Note: Stocks above the horizontal divider (Rank 1–22) belong to the high-coverage group (coverage >70%), and those below to the low-coverage group; data period January 1, 2013 to December 31, 2025, totaling 3,176 trading days.

Table A.1: News Coverage Details for 47 Constituent Stocks (2013–2025)

Rank	Stock	Coverage Rate	Trading Days with News	Total News Articles
1	2330 TSMC	100.0%	3,176	51,124
2	2317 Hon Hai	99.9%	3,172	36,611
3	2454 MediaTek	96.8%	3,069	19,426
4	2303 UMC	96.5%	3,059	16,963
5	2887 TS Financial	93.0%	2,948	12,436
6	1303 Nan Ya Plastics	93.0%	2,947	10,572
7	2891 CTBC Financial	92.8%	2,943	13,220
8	3008 Largan Precision	92.8%	2,943	23,775
9	2882 Cathay Financial	87.9%	2,786	15,046
10	2881 Fubon Financial	86.0%	2,727	12,984
11	2883 KGI Financial	85.2%	2,701	9,124
12	2002 China Steel	81.9%	2,596	9,310
13	2886 Mega Financial	80.9%	2,562	9,930
14	2890 SinoPac Financial	79.7%	2,527	6,648
15	3231 Wistron	77.5%	2,458	8,060
16	2382 Quanta Computer	76.4%	2,424	8,242
17	2308 Delta Electronics	75.3%	2,389	6,700
18	2885 Yuanta Financial	75.1%	2,382	7,046
19	2884 E.Sun Financial	74.9%	2,373	6,634
20	2412 Chunghwa Telecom	74.4%	2,358	6,273
21	1301 Formosa Plastics	73.4%	2,325	6,988
22	2357 ASUS	70.6%	2,238	6,668
23	2892 First Financial	68.1%	2,161	4,978
24	2603 Evergreen Marine	66.2%	2,099	6,853
25	3045 Taiwan Mobile	64.8%	2,054	4,399
26	2408 Nanya Technology	61.6%	1,946	5,551
27	2327 Yageo	60.3%	1,898	7,063
28	4904 Far EasTone	54.3%	1,721	3,686
29	3034 Novatek Microelectronics	51.8%	1,644	3,954
30	6505 Formosa Petrochemical	51.3%	1,621	3,529
31	2880 Hua Nan Financial	50.4%	1,596	2,746
32	2379 Realtek Semiconductor	49.7%	1,574	3,478
33	3661 Alchip Technologies	49.1%	1,445	3,442
34	2345 Accton Technology	47.4%	1,502	3,374
35	2615 Wan Hai Lines	47.0%	1,483	4,542
36	2301 Lite-On Technology	46.2%	1,463	2,927
37	2383 Elite Material	45.8%	1,452	3,028
38	3665 Bizlink	42.8%	1,344	2,704
39	3017 Asia Vital Components	42.6%	1,350	3,052
40	5880 Taiwan Cooperative Financial	40.3%	1,274	1,989
41	2395 Advantech	39.8%	1,261	2,200
42	2912 President Chain Store	37.5%	1,188	1,905
43	2207 Hotai Motor	34.3%	1,086	1,861
44	3653 Jentech Precision	28.1%	866	1,445
45	2360 Chroma ATE	19.5%	613	939
46	1216 Uni-President	19.0%	603	848
47	2059 King Slide Works	12.0%	381	549
0050 Market Index		99.1%	3,152	34,406