**Author**: Radoslav Geogiev

**Supervisor**: Dr. Fenghui Yu

# Multiple Pairs Trading for Portfolio Optimization with Reinforcement Learning

## 1. Background

- **Pairs trading** is a popular trading strategy reliant on the mean-reverting properties of two assets.
- **Portfolio optimization** deals with the dynamic reallocation of stock weights.
- **PPO** is an actor-critic on-policy **RL strategy** that supports a **continuous action space**.
- Knowledge gap for RL strategies for portfolio optimization with **multiple pairs**.

## 2. Objectives

Develop a **pair-unaware (Unstructured)** and a **pair-aware (Structured)** PPO models. Compare the two models to answer the following research questions:

- **RQ1**: **"**How does the Structured model for portfolio optimization with asset pairs compare to the Unstructured model?"
- **RQ2**: "How does the performance of the Structured and Unstructured models change in terms of realised profits with a different number of asset pairs?"

## 3. Methodology

- **Stocks** are selected from the energy sector of the **S&P 500** via Yahoo Finance. **Training period** is 01.01.2015-01.01.2022. **Testing period** is 01.01.2022-0.1.01.2024.
- **Pairs of stocks** are evaluated for **cointegration** using the **Johansen test**, and the top 10 with the highest trace statistic are used.
- The **Unstructured** model trades directly with the stocks and a risk-free bank account.
- The **Structured** model trades only on the spreads of the pairs and a risk-free bank account. Additionally, it keeps track of specific information regarding the pairs in its state.

## 4. Results

### 4.1 Portfolio Value Movements

Structured model and equal-weight benchmark maintain higher portfolio values than the Unstructured model throughout the whole testing period.

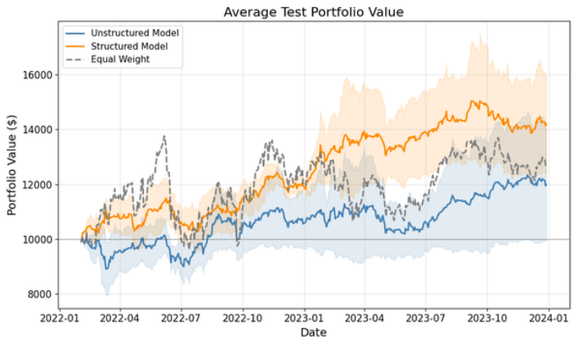


Figure 1: Average portfolio values of the Unstructured and Structured models across the testing period (01.01.2022 to 01.01.2024).

Portfolio value of the Unstructured model seems to be much more dependent on the overall market movement.

### 4.2 Performance Metrics

Overall, the Structured model achieves the highest results across all 5 metrics.

- The Structured model has a Sharpe ratio that is 0.773 higher than the Unstructured model, and 0.590 higher than the equal-weight benchmark.
- Annual return of the Structured model is 11.5% higher than that of the Unstructured model.
- The Structured model has lower volatility across all 5 metrics except for the maximal drawdown.

Table 1: Performance metrics of the Unstructured and Structured models. Standard deviation is shown in brackets.

| | Unstructured | Structured | EW |
|---|---|---|---|
| End Value ($) | 11,985 (2,052) | 14,198 (1,808) | 12,670 |
| Ann. Return (%) | 9.6 (9.7) | 20.1 (8.0) | 13.3 |
| Ann. Vol. (%) | 20.5 (7.9) | 18.0 (5.4) | 32.6 |
| Sharpe Ratio | 0.104 (0.458) | 0.877 (0.324) | 0.287 |
| Max Drawdown (%) | −20.2 (1.9) | −13.3 (2.6) | −30.1 |

### 4.3 Weight Allocation

- The Unstructured model prefers to invest in WMB, MPC, and FANG, and to short-sell BKR and DVN.
- Structured model assigns positive weights to all pairs, especially MPC/WMB.
- Both models are quite reliant on the risk-free bank account.
- There is a high volatility in the assignment of weights.
- Both models display a strong correlation between weight assignment and the returns of the stocks/pairs.

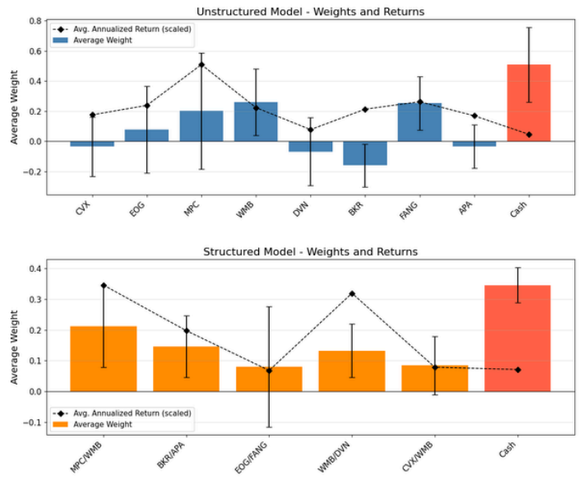


Figure 2: Average weights allocated to assets/spreads and risk-free bank account (cash). Standard deviation of the weights is included. The scaled annual return of the assets/spreads and cash is displayed with a black line.

### 4.4 Analysis across different number of pairs

- The Structured model obtains the highest annual returns and Sharpe ratios in the cases of 3, 5, and 10 pairs.
- The highest Sharpe ratio of 0.877 was obtained by the Structured model for the case of 5 pairs.
- The variance of the Unstructured model increases in the case of annual returns and decreases in the case of Sharpe ratios as we increase the number of pairs.
- No notable patterns between performance and the number of pairs are present in the Structured model.

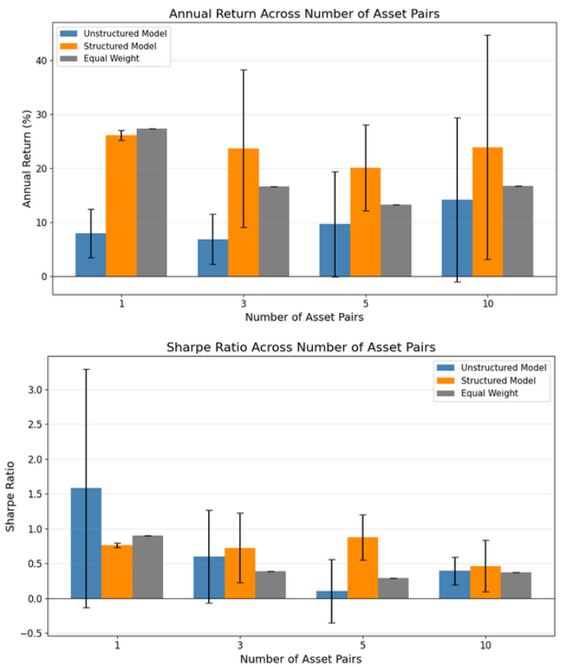


Figure 3: Annual returns and Sharpe ratios of the Unstructured and Structured models across different numbers of asset pairs. Equal-weight strategy values have been included. Standard deviations are also displayed.

## 5. Conclusions

- The added structure to the pair-aware model helps it achieve much higher results in terms of portfolio value, annual returns, and Sharpe ratio.
- The Structured model manages to consistently outperform the Unstructured model and the equal-weight benchmark across a different number of asset pairs.
- No notable relationship between performance and the number of pairs was discovered.
- The results showcase that a multi-pair RL model approach to portfolio optimization might be promising and emphasize the need for further research.