Application of Operations Research Techniques in Trading

A Comparative Study of Optimization Methods in Pair-Based Trading and Technical-Indicator Strategies

[1]Arin Singh, [2]Harish Patil, [3]Pragatheeswaran Raja, [4]Rishabh Sharma, [5]Granth Patel

[1-5]Department of B.Sc. Finance

[1-5]NMIMS (Deemed-to-be University), Bengaluru

***Abstract*—Trading involves a chain of decisions: which assets to hold, in what proportion, when to enter and exit, how much risk to accept, and which parameters should drive a rule-based signal. Operations Research (OR) treats each of these as a formal optimization problem with decision variables, constraints and an objective. This paper presents a comparative review of six studies that apply different OR methods to trading. Three address pair-based (statistical-arbitrage) trading: Majorization-Minimization for nonconvex mean-reverting portfolio design, optimal stopping for entry and exit timing, and bi-objective optimization for balancing profit against risk. The other three optimize the parameters of widely used technical indicators: exhaustive grid search for the MACD, a genetic algorithm for Bollinger Bands, and a two-stage genetic algorithm and artificial neural network for the Relative Strength Index (RSI). The comparison is qualitative, because the studies use different markets, periods, constraints and performance measures. The review finds that the method must be matched to the structure of the decision: nonconvex programming suits static allocation, stochastic control suits timing under uncertainty, multi-objective methods suit conflicting goals, and exhaustive or evolutionary search suits low-dimensional parameter tuning. Every method depends on assumptions that may not hold out of sample, so rigorous out-of-sample validation and explicit treatment of transaction costs are essential.**

***Index Terms*—Operations Research, pair trading, statistical arbitrage, Majorization-Minimization, optimal stopping, multi-objective optimization, genetic algorithm, grid search, MACD, Bollinger Bands, RSI, neural networks**

I. Introduction

A trader rarely faces a single decision. A market-neutral strategy, for example, must first decide which assets to combine and with what weights, then decide when to open and close the position as prices drift, and finally decide how much risk is acceptable for the return sought. A directional trader using a technical indicator faces a different but related decision: the indicator's parameters, such as look-back windows and thresholds, determine every buy and sell signal, yet they are usually left at conventional default values. In each case, intuition and rules of thumb are unreliable because relationships between assets are estimated from noisy data and change over time.

Operations Research offers a disciplined alternative. A trading decision is expressed through decision variables (weights, entry and exit points, indicator parameters), constraints (a fixed budget, a variance floor, admissible parameter ranges) and an objective (mean reversion, expected payoff, risk-adjusted return). Once stated in this form, the problem can be solved with a method suited to its mathematical structure: nonconvex optimization, stochastic control, multi-objective programming, exhaustive enumeration, or evolutionary and learning-based search.

This paper reviews six studies that illustrate this range. Three concern pair-based trading, a strategy that trades the relative price movement of related assets rather than the direction of any one of them (Zhao &

Palomar, 2017; Song & Zhang, 2021; Yang & Malik, 2024). The remaining three concern technical-indicator trading and tune the parameters of the MACD (Kang, 2021), Bollinger Bands (Siquieri & Costa, 2024) and the RSI. Together they cover the three stages of a trading system that OR can improve: construction, timing and calibration.

The review compares the studies on their modelling logic, solution method, reported benefits and limitations. It does not reproduce their numerical experiments or claim a new empirical ranking, since that would require applying every method to a common dataset under consistent assumptions.

II. REVIEW SCOPE AND METHODOLOGY

The review is based on six studies supplied for the group project, organized in two sets. The first set comprises three papers on pair-based trading, each addressing a different stage of the strategy. The second set comprises three papers on parameter optimization of technical indicators, each using a different search technique. Grouping the studies by decision problem and solution method, rather than describing them in isolation, makes it possible to see how the choice of OR technique follows from the structure of the problem.

A comparative reading framework was applied to each study: (i) the trading problem being addressed; (ii) the decision variables and objective; (iii) the OR method or search algorithm used; (iv) the benefits reported; and (v) the limitations that can reasonably be inferred. The synthesis is qualitative. The papers use different assets (futures, individual stocks, multi-asset spreads), different sample periods and different performance indicators, so their numerical results are not directly interchangeable.

No new price data were downloaded and no trading experiment was run. The Results and Discussion section therefore reports patterns in the reviewed literature rather than newly calculated investment performance. Figures quoted from the original studies are attributed to them and should be read in the context of their own data and assumptions.

III. RESEARCH OBJECTIVES

a) To identify the principal OR methods used in the selected trading studies and the decision each one optimizes.
b) To compare how the studies formulate objectives, constraints and risk, and how those choices shape the resulting strategy.
c) To assess the advantages and limitations of exact optimization, exhaustive search, metaheuristics and learning-based approaches in a trading context.
d) To connect the literature with core OR concepts, including nonconvex and multi-objective optimization, stochastic control, heuristic search and sensitivity analysis.
e) To identify a practical design for future research that compares methods under consistent data and cost assumptions.

IV. LITERATURE REVIEW

4.1 Portfolio construction for mean reversion: nonconvex optimization

Pair-based trading is a market-neutral strategy. Rather than betting on the direction of one asset, the trader combines related assets so that the resulting price spread fluctuates around a stable mean, then trades deviations from that mean in the expectation that they will close. The first decision is how to build the spread. Zhao and Palomar (2017) treat this as the design of a portfolio of several correlated assets whose spread reverts to its mean as strongly as possible. The difficulty is that mean-reversion strength, portfolio variance and the budget all depend jointly on the same weight vector, so improving one property can worsen another.

The authors model the problem as a constrained nonconvex optimization: a mean-reversion criterion is optimized subject to a variance floor, which prevents the portfolio from collapsing into a flat and untradeable line, and a budget constraint fixing the capital invested. Because the objective is a nonconvex

function of the weights, standard convex solvers cannot be applied directly. They therefore use Majorization-Minimization (MM). At each iteration a simpler convex surrogate that bounds the true objective from above is minimized, through a sequence of generalized trust-region subproblems. The solutions improve the original objective monotonically until the algorithm converges to a stationary point. In the reported backtests the MM-designed portfolio outperformed each individual underlying spread and a prior benchmark method.

The principal benefit is that portfolio construction becomes a repeatable, mathematically specified procedure that accommodates explicit constraints, removing much of the subjectivity involved in choosing assets and weights by inspecting correlated charts. The limitations are equally clear. MM guarantees convergence only to a stationary point, not a global optimum, so the solution can depend on initialization. The method also assumes that estimated mean reversion will persist, and a portfolio that is optimal on historical data can underperform live because transaction costs and regime shifts are absent from the objective.

4.2 Trade timing: optimal stopping

Once a spread exists, the trader must decide when to open and close a position. A fixed rule, such as trading whenever the spread is two standard deviations from its mean, ignores how the spread is behaving at that moment and may enter too early or exit too late. Song and Zhang (2021) recast the choice as an optimal stopping problem, a branch of stochastic control in which a decision-maker observes an uncertain process over time and chooses the moment to act that maximizes the expected outcome, weighing the value of acting now against the value of waiting for a better state.

In this framework the spread is modelled as a stochastic process. At each point in time the trader observes its state and decides whether to trade immediately or continue waiting. The quantity optimized is the expected trading payoff, and the solution is a stopping rule, typically boundaries on the spread's value or on time remaining, that states exactly when to act. The rule is then evaluated against fixed-threshold alternatives. The improvement over arbitrary cutoffs comes from deriving the boundaries from the spread's own dynamics rather than choosing them by convention, and from treating trading as a sequence of decisions under uncertainty rather than a single guess.

The weakness is model dependence. The stopping rule is optimal only under the assumed stochastic model of the spread. If the true dynamics differ, or the relationship between the paired assets shifts, the optimality guarantee no longer holds. Transaction costs, slippage and other frictions are often simplified in such models, and including them can change the optimal boundaries substantially.

4.3 Balancing return and risk: bi-objective optimization

A strategy chosen only to maximize expected profit can carry unacceptable risk, while one chosen only to minimize risk may earn too little to be worthwhile. Yang and Malik (2024) formulate the trader's problem as a bi-objective optimization: expected profit is maximized while risk is simultaneously controlled. The objectives conflict because strategies that chase higher returns typically need larger and more concentrated positions, which raises variance and drawdown.

Instead of merging the two goals into one objective through an arbitrary weighting, the model is solved for a set of Pareto-optimal solutions, points at which profit cannot be increased without increasing risk. These points form an efficient frontier from which the trader selects the strategy matching a chosen risk appetite, and the resulting strategies are backtested for realized return and risk. The contribution is to avoid a premature decision about how much risk to accept: the trade-off is made explicit and the trader is offered a menu of efficient strategies rather than a single optimum for one possibly mismatched criterion.

The limitation is that the frontier depends heavily on how profit and risk are each defined, and different traders will reasonably prefer different points on it. The method therefore does not remove the need for a judgment about risk preference. Parameter choices can alter which strategies appear efficient, and historical performance along the frontier does not guarantee future performance.

4.4 Calibrating the MACD: exhaustive grid search

The remaining three studies move from pair construction to the calibration of indicator-based rules. Kang (2021) examines the Moving Average Convergence Divergence (MACD) indicator, which is almost always applied with the conventional (12, 26, 9) setting. The problem is that these values are not necessarily optimal: they differ from the two separate parameter sets Appel, the indicator's developer, suggested for buying and selling, and they were derived for a six-day trading week rather than the modern five-day week.

The study uses exhaustive grid search, a brute-force method that evaluates every combination in a defined range. With the fast, slow and signal periods ranging over 3–20, 5–40 and 3–40, the grid contains 24,624 combinations. Of these, 5,168 are infeasible because the slow moving average period is shorter than the fast period, leaving 19,456 valid models. Each was tested against the traditional (12, 26, 9) setting on Nikkei 225 futures over 2011–2019. Three supplementary rules, namely a minimum holding period, a percentage-threshold histogram crossover and a first-peak or first-bottom search, were also tested as ways to filter false signals.

The traditional setting produced a net loss over the period, while an example optimized model, MACD (4, 22, 3), produced a substantial net gain. The supplementary rules meaningfully improved optimized models but did not rescue the traditional one, which supports the paper's conclusion that optimizing the parameters themselves should take priority over layering on additional filters. The limitations are computational cost and narrow validation: the search is thorough but expensive, and it was validated on a single market and period. Because it selects the best of many thousands of in-sample results, it is also exposed to the risk that the winning combination reflects chance.

4.5 Bollinger Bands: genetic algorithm

Siquieri and Costa (2024) address the default Bollinger Bands settings of a 20-period average and a band width of two standard deviations, which need not suit every stock. Because each candidate parameter set must be backtested to be evaluated, the objective is costly and irregular, and a genetic algorithm (GA) is used to search for better values. The GA optimizes three parameters: the averaging period N (15–40), the band-width multiplier K (1–4, a continuous variable) and the type of moving average (simple or exponential). The trading rule is deliberately simple: buy when the price crosses back above the lower band and sell when it crosses back above the central band. The fitness function is the return of the trading system.

The GA used a population of 8 solutions per generation over 100 generations, steady-state selection, single-point crossover and a 30% mutation probability. It was applied to one year of intraday minute data for three technology stocks (GOOG, MSFT and AAPL), with each optimization using a three-week training window and the following one-week test window. The optimized parameters outperformed the standard parameters in about 77.8% of the 54 test periods, a clear and consistent edge. However, the size of the improvement was modest, roughly 1–4%, so the practical gain per period was small despite the high win rate. The design also shows the value of the rolling train-and-test structure, since it evaluates parameters on data that were not used to find them.

4.6 RSI: genetic algorithm combined with a neural network

The final study applies the same logic to the Relative Strength Index, whose conventional settings are a 14-day period with 30 and 70 thresholds. Its aim goes beyond tailoring parameters to each stock: it seeks to predict good parameters for future periods rather than only optimizing on history. The method has two stages. In the first, a GA finds the best RSI parameters for each stock, using a population of 1,000 over 100 generations. Each chromosome has three genes, the upper threshold (60–90), the lower threshold (10–40) and the number of observations (12–28), and the GA uses tournament selection, two-point crossover and bit-flip mutation with probability 0.05.

In the second stage, the GA output trains an artificial neural network to predict the optimal parameters for a period the GA has not analyzed. The network has an input layer of 445 nodes (RSI buy and sell signals), one hidden layer of 10 neurons with ReLU activation, and three linear output nodes, one for each RSI parameter. It was trained with the Adam optimizer and mean squared error loss over 2,000 epochs on an 80/20 training and test split. Of the network's predictions, 58.51% both beat the traditional RSI returns and were positive. Treated as an equal-weighted portfolio across 94 stocks, the AI-driven approach returned

22.71% annually against 4.55% for the fixed-parameter approach. Just under half of the individual predictions did not outperform the traditional method, however, so the results are strong in aggregate but not reliable stock by stock, and the approach requires considerably more modelling and computational complexity than the fixed-parameter benchmark.

V. COMPARATIVE ANALYSIS AND DISCUSSION

5.1 From trading problems to OR models

All six studies convert an informal trading idea into a formal decision problem. Weights, entry and exit boundaries and indicator parameters become decision variables; a budget, a variance floor or a parameter range become constraints; and mean reversion, expected payoff, risk or backtested return become objectives. The solution technique then follows from the structure of that problem. A static allocation with a smooth but nonconvex objective calls for nonconvex programming with MM. A timing decision under a stochastic process calls for stochastic control. Two conflicting goals call for Pareto analysis. A small, bounded parameter space can be enumerated, while a larger or irregular space with a costly, simulation-based objective suits evolutionary search, optionally supported by a learning model.

5.2 Comparative summary of the main approaches

Study	OR method	Decision optimized	Main strength	Main limitation
Zhao & Palomar (2017)	Nonconvex optimization with Majorization-Minimization	Asset weights of the mean-reverting spread	Repeatable procedure with explicit variance and budget constraints	Only a stationary point is guaranteed; costs and regime shifts excluded
Song & Zhang (2021)	Optimal stopping (stochastic control)	When to enter and exit	State-dependent rule replaces arbitrary thresholds	Optimal only under the assumed spread model; frictions simplified
Yang & Malik (2024)	Bi-objective optimization (Pareto frontier)	Balance of profit and risk	Makes the trade-off explicit; tunable to risk appetite	Frontier depends on how profit and risk are defined
Kang (2021) – MACD	Exhaustive grid search	Three MACD parameters	Thorough; guaranteed to find the best in-sample combination	Computationally expensive; one market and period; selection bias
Siquieri & Costa (2024) – Bollinger Bands	Genetic algorithm	N, K and moving-average type	Efficient search of a costly backtest objective; rolling test windows	Modest gains (about 1–4%); no global optimum guarantee
RSI study – AI	Genetic algorithm plus neural network	RSI thresholds and period, predicted for new periods	Forward-looking; strong portfolio-level return	Inconsistent stock by stock; high complexity

Table 1. *Qualitative comparison synthesized from the reviewed studies. The table is not a ranking and does not imply that all methods were tested on a common dataset.*

5.3 Exact optimization, enumeration and heuristic search

The studies differ in the guarantees they offer. The MM algorithm and the optimal stopping rule come with mathematical structure: MM improves the objective monotonically and converges to a stationary point, and

the stopping boundaries are optimal for the assumed process. Grid search offers a different guarantee, since it examines every candidate and so cannot miss the best in-sample combination, but it becomes impractical as the number of parameters or the width of the ranges grows. The genetic algorithm trades certainty for scalability: it samples the search space intelligently and can handle continuous and categorical parameters together, as in the Bollinger study, but it cannot prove that no better solution exists, and its result depends on population size, mutation rate and random initialization. The neural network in the RSI study adds a prediction layer on top of search, which extends the idea from describing the past to anticipating the future but introduces its own training and generalization risks.

5.4 The common weakness: assumptions and overfitting

A shared limitation runs through all six studies. The pair-trading models rely on a spread that keeps reverting and, in the case of optimal stopping, on a particular stochastic model. The indicator studies select parameters that performed best on history. In both cases an optimum found on past data is not guaranteed to hold in the future, and the more candidates that are searched, the greater the chance that the winner reflects noise. Transaction costs and slippage, which erode the small edges reported in some of these studies, are absent or simplified in most of the objectives. This is why the rolling train-and-test windows in the Bollinger study and the held-out test split in the RSI study are as important as the optimizer itself.

5.5 Connection with Operations Research coursework

The reviewed papers map directly onto core OR ideas: decision variables, objective functions and constraints; nonconvex and convex optimization and the use of a surrogate problem to solve a hard one; optimization under uncertainty and sequential decision-making; multi-objective trade-offs and Pareto efficiency; heuristic and metaheuristic search; and sensitivity analysis, which asks how far a solution changes when the data or assumptions change. The techniques also generalize well beyond trading: nonconvex optimization to engineering design, optimal stopping to equipment replacement and real options, multi-objective optimization to supply-chain and production planning, and genetic algorithms to scheduling and routing.

VI. RESULTS AND DISCUSSION

Table 2 collects the principal findings reported by the original authors. These are reproduced from the reviewed studies and are not the result of any new calculation; because the studies use different markets, periods and metrics, the figures should not be compared across rows.

Study	Data and setting	Principal reported finding
Zhao & Palomar (2017)	Backtests of designed spreads against single-asset spreads and a prior benchmark	MM-designed portfolio outperformed each individual spread and the benchmark method
Song & Zhang (2021)	Optimal stopping rule evaluated against fixed-threshold rules	State-dependent boundaries preferred to a fixed threshold under the assumed model
Yang & Malik (2024)	Pareto-optimal strategies backtested for return and risk	Efficient frontier exposes the profit–risk trade-off for different risk appetites
Kang (2021) – MACD	Nikkei 225 futures, 2011–2019; 19,456 valid parameter sets	Traditional (12, 26, 9) setting lost money; optimized (4, 22, 3) produced a substantial gain
Siquieri & Costa (2024) – Bollinger	GOOG, MSFT, AAPL minute data; 54 test periods	GA parameters beat standard settings in about 77.8% of periods; average gain about 1–4%

Study	Data and setting	Principal reported finding
RSI study – AI	94 stocks; GA then neural network, 80/20 split	58.51% of predictions beat traditional RSI and were positive; 22.71% vs 4.55% annual portfolio return

Table 2. *Principal findings as reported in the reviewed studies.*

Several patterns emerge. First, default parameters are rarely optimal. The MACD and Bollinger results both show that conventional settings can be improved, and in the MACD case the standard setting produced a loss over the sample period. Second, the method should match the decision. Pair-trading research shows that construction, timing and risk balance are distinct problems that need distinct tools, and no single technique covers all three. Third, the size of the gain varies widely. The Bollinger improvement was modest per period, whereas the RSI portfolio result was large, yet that result also came with the greatest modelling complexity and with inconsistent performance on individual stocks.

Fourth, the order of priorities matters. The MACD study found that optimizing the core parameters should take precedence over adding filters, which suggests that calibrating the main model usually yields more than patching a poorly calibrated one. Finally, practical value depends on evidence beyond the optimizer. Reported advantages are strongest where the study separates the data used to find parameters from the data used to judge them, and weakest where validation covers one market, one period or a simplified cost structure.

VII. RESEARCH GAPS AND FUTURE SCOPE

Three gaps stand out. The first is the absence of a common benchmark: because each study uses its own data and metrics, the methods cannot be ranked, and a controlled comparison on the same assets, periods, cost assumptions and performance measures is needed. The second is the treatment of frictions. Transaction costs, slippage and liquidity are absent or simplified in most of the objectives, yet they can erase the small edges reported and can shift optimal boundaries, as noted for optimal stopping. The third is robustness to regime change. All the methods assume that estimated relationships persist, so future work could incorporate regime detection, robust or adaptive formulations, and walk-forward testing across several markets.

Promising extensions follow from the reviewed work. The pair-trading models could be combined, for example by using an MM-designed spread, an optimal stopping rule for timing and a Pareto analysis of the resulting strategies within one framework. Indicator parameters could be optimized for risk-adjusted rather than raw return, using multi-objective evolutionary algorithms. Finally, penalties for overfitting, such as multiple-testing adjustments to exhaustive search, would make the reported optima more credible.

VIII. CONCLUSION

The reviewed studies show that Operations Research does not offer one universal technique for trading but a family of frameworks, each matched to a different decision. Nonconvex optimization with Majorization-Minimization supports portfolio construction, optimal stopping supports entry and exit timing, and bi-objective optimization makes the balance between return and risk explicit. For indicator-based trading, exhaustive grid search, genetic algorithms and a hybrid genetic-algorithm and neural-network approach each show that default parameters can be improved by systematic search, with the right technique depending on the size and structure of the search space.

The review does not identify a universally superior method, because the studies do not share a common empirical design and their objectives differ. The strongest practical lesson is to match the model to the decision and to treat optimization output as a hypothesis to be validated rather than a guaranteed result. Each method's usefulness ultimately depends on the validity of its assumptions and on the data used to estimate them. A rigorous future comparison would use identical assets, windows, costs and benchmarks with out-of-sample testing and sensitivity analysis. In this sense OR is most valuable not as a search for one best

algorithm, but as a disciplined way to make trading decisions explicit, testable and better aligned with the risks traders actually face.

References

[1] Kang, B.-K. (2021). Improving MACD technical analysis by optimizing parameters and modifying trading rules: Evidence from the Japanese Nikkei 225 futures market. Journal of Risk and Financial Management, 14(1), 37.

[2] [Author(s)]. [Year]. Artificial intelligence in stock market investment through the RSI indicator. [Source details to be completed from the original paper.]

[3] Siquieri, W., & Costa, O. L. V. (2024). Bollinger Bands trading strategy optimization with genetic algorithm. Congresso Brasileiro de Automática (CBA 2024).

[4] Song, S., & Zhang, H. (2021). Optimal stopping approach to pairs trading.

[5] Yang, Y., & Malik, A. (2024). Bi-objective optimization for pairs trading strategies.

[6] Zhao, Z., & Palomar, D. P. (2017). Mean-reverting portfolio design with budget constraint. arXiv:1701.05016.