

APPLICATION OF OPERATIONS RESEARCH TECHNIQUES IN PORTFOLIO OPTIMIZATION

A Comparative Study of Risk–Return Models

¹Vinit Sharma, ²Anvita Shetty, ³Sonam Goradia, ⁴Nandakishor Harilal, ⁵Ishita Kulkarni,
⁶Tejash Choudhary

¹⁻⁶B.Sc. Finance

¹⁻⁶NMIMS (Deemed-to-be University), Bengaluru

Abstract—Portfolio optimization is an allocation problem in which an investor seeks an appropriate balance between return, risk and practical investment restrictions. Operations Research (OR) provides a formal framework for this decision by defining decision variables, objective functions, constraints and solution procedures. This paper presents a comparative review of 17 distinct studies represented in the 18 supplied PDF files, two files contain the same article on pentagonal neutrosophic linear programming and are treated as one source. The reviewed literature covers mean–variance and mixed-integer quadratic programming, Value-at-Risk (VaR), Conditional Value-at-Risk (CVaR), linear-programming formulations, robust and multistage stochastic optimization, fuzzy and neutrosophic models, data envelopment analysis, Monte Carlo simulation, evolutionary and swarm-based metaheuristics, and a neurodynamic approach based on spiking neural networks. The comparison is qualitative rather than a pooled empirical test because the studies use different data, constraints, assumptions and performance measures. The review finds that model choice depends on the type of risk an investor wants to control, the reliability of estimated inputs, the restrictions imposed on holdings and the computational resources available. Classical mathematical programming offers transparency and, in many formulations, strong solution guarantees, metaheuristics are useful for difficult discrete search spaces but do not automatically guarantee a global optimum. Uncertainty-aware methods broaden the model but require additional assumptions and calibration. The paper concludes that no single method is universally best and that meaningful comparison requires a common out-of-sample testing framework.

Index Terms—Operations Research, portfolio optimization, mean–variance, VaR, CVaR, robust optimization, stochastic programming, fuzzy returns, metaheuristics, portfolio risk

I. Introduction

Portfolio optimization asks how an investor should divide available capital among a set of financial assets. The decision is not simply to choose the asset with the highest expected return. Returns are uncertain, assets may move together, and an allocation that looks attractive on its own can increase the risk of the portfolio as a whole. A useful portfolio therefore depends on both the expected behaviour of individual assets and the relationships among them.

Operations Research makes this decision more structured. Portfolio weights become decision variables; expected return and risk become objective functions or constraints; and investment rules are expressed mathematically. Typical rules include investing the full budget, prohibiting short sales, limiting exposure to a single asset or sector, selecting no more than a fixed number of securities, and accounting for transaction costs. Once expressed this way, the portfolio problem can be solved through linear programming (LP),

quadratic programming (QP), mixed-integer programming, simulation, stochastic optimization or heuristic search.

The classical mean–variance framework associated with Markowitz remains an important starting point because it formalizes the trade-off between expected return and variance. Yet variance treats positive and negative deviations around the mean symmetrically. Investors who are particularly concerned about severe losses may prefer downside-oriented measures such as VaR or CVaR. Other investors may focus on risk contribution, utility, robust performance under uncertain inputs, or the possibility that returns cannot be represented reliably by a single probability distribution.

The selected papers reflect this variety. Some develop mathematical formulations that can be solved by established optimization algorithms. Others adapt nature-inspired search procedures to portfolios with cardinality constraints or large numbers of assets. A further group considers fuzzy or neutrosophic representations of uncertain returns, while multistage stochastic models examine decisions that unfold over time. These studies are related by their use of formal decision models, but they do not all solve the same problem.

This review compares the methods on their modelling logic, risk representation, computational characteristics and practical limitations. It does not claim to reproduce the numerical results of the original papers or to establish a new empirical ranking. Such a ranking would require applying the models to a common dataset with consistent assumptions.

II. Review Scope and Methodology

The review uses the 18 PDF files supplied for the project as its primary source base. The files represent 17 distinct articles because the paper by Khalifa and co-authors on pentagonal neutrosophic linear programming was uploaded twice. The duplicate is not counted as a separate study. The sources span classical risk measures, linear and quadratic programming, robust and stochastic models, fuzzy portfolio selection, DEA, Monte Carlo methods and metaheuristic algorithms. Duplicate copies are counted once; the bibliography lists the distinct studies used in the synthesis.

A comparative reading framework was used. For each distinct study, attention was given to: (i) the portfolio problem being addressed; (ii) the risk or performance measure; (iii) the optimization formulation or search algorithm; (iv) the type of constraints or uncertainty considered; and (v) the contribution and limitations that can reasonably be inferred from the paper. The review then groups the studies by methodological family rather than describing them in a disconnected sequence.

The synthesis is qualitative. The papers use different markets, asset universes, sample periods, return assumptions, constraints and performance indicators. Their numerical results are therefore not directly interchangeable. A result obtained on a particular exchange or benchmark portfolio does not establish that the same algorithm will outperform in another market. Similarly, a method designed for a discrete portfolio with a fixed number of holdings should not be compared without qualification to a continuous-weight model with no cardinality restriction.

No new return series were downloaded and no portfolio experiment was run for this review. Accordingly, the Results and Discussion section reports patterns in the selected literature, not newly calculated investment performance. This distinction avoids presenting a literature synthesis as if it were an original empirical test.

III. Research Objectives

1. To identify the principal Operations Research methods used in the selected portfolio-optimization studies.
2. To compare the risk measures and modelling assumptions that shape the resulting allocation decisions.
3. To assess the advantages and limitations of exact mathematical programming, uncertainty-aware formulations and metaheuristic search.
4. To connect the literature with core OR concepts, including LP, QP, mixed-integer programming, simulation, heuristic search and sensitivity analysis.
5. To identify a practical design for future research that could compare methods under consistent data and trading assumptions.

IV. Literature Review

4.1 Classical risk–return models and downside risk

Markowitz's mean–variance approach provides the conceptual baseline for much of portfolio theory. In a standard formulation, the investor minimizes portfolio variance for a required expected return, or maximizes expected return after penalizing variance. Portfolio variance depends on the covariance matrix, so diversification is represented directly: assets that do not move in perfect synchrony can reduce overall portfolio risk. When the covariance matrix is suitable and the constraints are convex, the problem can often be formulated as a quadratic program. (Markowitz, 1952)

The strength of this framework is its clarity. The model makes the return–risk trade-off explicit and gives analysts a way to trace an efficient frontier. Its weaknesses are also well known in the context of the reviewed studies. Variance penalizes upside and downside deviations, estimates of expected returns can be unstable, and the optimized weights may be sensitive to small changes in input data. Realistic restrictions—especially a fixed number of holdings or transaction costs—can also turn a relatively tractable continuous problem into a harder mixed-integer one.

Rockafellar and Uryasev's work on CVaR is central to the downside-risk strand of the literature. VaR describes a loss threshold at a chosen confidence level, but by itself it does not describe how severe losses can be beyond that threshold. CVaR, often interpreted as expected tail loss under suitable conditions, focuses on losses in the tail and has useful coherence properties. Its key OR contribution is a formulation that can be expressed using an auxiliary threshold variable and scenario-based excess-loss variables, enabling linear-programming techniques in important settings. (Rockafellar & Uryasev, 2002)

Palczewski's PortfolioOptim package demonstrates how several downside-risk measures—including CVaR, mean absolute deviation (MAD), lower semi-absolute deviation (LSAD) and dispersion CVaR—can be handled through LP algorithms for discrete return distributions. The paper is important not only for its risk measures but also for translating formulations into computational procedures. The package-oriented contribution illustrates the value of reproducibility and implementation: a model is more useful when the algorithm and assumptions can be applied in a transparent way. (Palczewski, 2018)

Cesarone, Martino and Tardella extend the mean–variance framework by incorporating VaR and formulate the resulting model as a mixed-integer quadratic program. Their work highlights a tension between the popularity of VaR in risk management and the computational difficulty of optimizing it directly. The inclusion of VaR can better reflect a particular risk threshold, but the resulting formulation may require more complex solution procedures than a basic convex QP. (Cesarone et al., 2023)

Geissel and co-authors examine optimal expected utility risk measures in a risk-constrained portfolio context and compare them with VaR-based constraints. Their approach permits the investor's risk attitude and time preference to influence the risk specification. This is conceptually useful because risk is not identical for every investor: a loss that is acceptable to one investor may be unacceptable to another. At the same time, utility-based models depend on the chosen utility specification and its parameters, so their outputs must be interpreted in light of those choices. (Geissel et al., 2022) Their rolling-forecast case study using three indices reports that both the OEU-constrained and VaR-constrained strategies performed considerably better than equal-weight and buy-and-hold benchmarks, while the OEU formulation allows investor-specific risk attitudes and time preferences (Geissel et al., 2022).

4.2 Robust and multistage stochastic optimization

Robust optimization responds to uncertainty in model inputs by considering a set of plausible parameter values rather than assuming that one estimate is exact. Georgantas, Doumpos and Zopounidis compare robust portfolio-selection approaches based on different risk measures. The significance of this work is the comparative perspective: robust optimization is not a single model, and the way uncertainty is represented can materially change the allocation. A robust solution may sacrifice some performance under the most optimistic estimate in exchange for greater stability across plausible scenarios. (Georgantas et al., 2021) Their empirical comparison uses US-market data from 2005–2020 and considers robust mean–variance, CVaR and Omega-ratio models; the key value for this review is that alternative robust formulations are assessed on out-of-sample performance rather than assumed to be interchangeable (Georgantas et al., 2021).

The practical challenge is selecting the uncertainty set. If it is too narrow, the model may not protect against meaningful estimation error. If it is too broad, the solution may become excessively conservative. Robustness therefore involves a modelling choice rather than a free improvement. Analysts should explain which parameters are uncertain, how the uncertainty range is calibrated and how the chosen range affects the resulting portfolio.

Gatzert and co-authors consider irreversible long-term renewable-energy investments under policy risk through a mixed-integer multistage stochastic model and a moving-horizon approach. Although this is not a conventional portfolio of listed shares, it extends portfolio optimization to investment opportunities whose decisions persist over time and whose future conditions depend on uncertain policy developments. The paper illustrates why a one-period model may be inadequate when investments cannot easily be reversed. Multistage stochastic programming represents decisions at different points in time and allows later actions to depend on information revealed in the meantime. (Gatzert et al., 2021)

The cost of this realism is computational and informational. Multistage models can grow quickly as the number of periods and scenarios increases. They also require assumptions about how uncertainty evolves and which decisions can be changed after new information arrives. For a student-level or practical portfolio study, the lesson is not that every model should be multistage; rather, the model's time structure should match the decision being studied.

4.3 Fuzzy, possibilistic and neutrosophic models

Probability-based models assume that uncertain returns can be represented using probabilistic information. In practice, expected returns and risk estimates may be imprecise, particularly when historical observations are limited or when an analyst relies on qualitative forecasts. Fuzzy and possibilistic approaches represent imprecision using fuzzy numbers and related measures rather than requiring every input to be specified as a conventional probability distribution.

Bhattacharyya develops portfolio-selection models based on a possibilistic Sharpe ratio. The approach defines possibilistic quantities such as mean, variance, standard deviation and risk premium for fuzzy returns, then uses the ratio of risk premium to dispersion as a risk-adjusted performance measure. This provides an alternative way to express risk aversion and return when inputs are fuzzy. However, fuzzy representations require the analyst to choose membership functions and interpret their parameters. Results are not automatically comparable with those from probabilistic Sharpe ratios unless the underlying assumptions are aligned. (Bhattacharyya, 2013)

Zhou and co-authors study portfolio performance with fuzzy returns using data envelopment analysis. Their paper questions the common practice of treating risk as an input and return as an output, arguing instead that initial wealth is the input and terminal wealth is the output, while risk and return are characteristics derived from terminal wealth. The use of nonlinear DEA and fuzzy returns broadens the idea of portfolio evaluation beyond a single scalar objective. It is useful when the purpose is to evaluate relative efficiency, although the result depends on the selected inputs, outputs and comparison set. (Zhou et al., 2019)

Khalifa and co-authors apply pentagonal neutrosophic linear programming to stock portfolio optimization. Neutrosophic representations seek to describe uncertainty using more than one dimension, including truth, indeterminacy and falsity components. The study illustrates an attempt to embed richer forms of uncertain information within an LP framework. Such models may be valuable where the uncertainty cannot be described adequately by a crisp point estimate, but the interpretation of the neutrosophic parameters and the conversion into an optimization model must be transparent for the results to be actionable. (Khalifa et al., 2022)

Taken together, these studies show that uncertainty can be represented in different ways. The choice between probability, fuzzy sets, possibility theory and neutrosophic numbers should be based on the information available and the question being asked. More elaborate uncertainty notation does not by itself guarantee more accurate portfolios; the assumptions must be defensible and the model must be tested against relevant outcomes.

4.4 Metaheuristics and hybrid optimization

Portfolio problems become harder when practical rules introduce discrete decisions. A cardinality constraint, for example, limits the number of assets that may receive a positive allocation. Minimum investment sizes, transaction costs and buy-in thresholds can also create non-convex or mixed-integer structures. In such cases, exact optimization may be computationally expensive, especially for large asset universes.

Evolutionary and swarm-based algorithms search for good solutions through repeated candidate generation and evaluation. Hochreiter proposes an evolutionary approach to risk-parity portfolio selection, combining a genetic algorithm with local search. Risk parity differs from simply minimizing variance: it seeks a portfolio in which assets make balanced contributions to total risk. The evolutionary framework is motivated by cases where long-short restrictions make the problem non-trivial. Its value is the ability to search a difficult solution space, although results depend on algorithm settings and do not automatically prove global optimality. (Hochreiter, 2015)

Liu and Xiao compare optimization algorithms in financial portfolio problems, including a genetic-algorithm approach and quadratic programming. Their study represents the broader question of whether search-based methods can improve practical results under a particular setup. Such comparisons are useful, but the reported ranking should be read in context: an algorithm may perform well under one dataset, objective and constraint set without being best under another. (Liu & Xiao, 2021)

Wijaya and Sidarto apply modified biogeography-based optimization to large-scale constrained portfolio optimization. The method is motivated by portfolio models whose dimensionality increases with the number of assets and constraints. Biogeography-based optimization uses the exchange and modification of candidate-solution characteristics to explore the search space. The paper contributes to the literature on applying nature-inspired methods to large constrained problems, while leaving the general caveat that stochastic search performance can vary with parameter choices and random initialization. (Wijaya & Sidarto, 2023)

The artificial bee colony study by Mansourinia and Momeni addresses cardinality-constrained portfolio optimization using a Tehran Stock Exchange case study. Cardinality restrictions are important because investors may prefer a manageable number of holdings, but they change the mathematical structure of the classical mean–variance problem. Bee-colony algorithms can explore candidate portfolios without requiring the entire problem to be transformed into a simple convex formulation. Their limitations include sensitivity to algorithm design, runtime and the need for fair comparisons with exact or other heuristic methods. (Mansourinia & Momeni, n.d.)

The simulated annealing–tabu search paper develops a hybrid approach for downside-risk portfolio optimization with cardinality constraints. Simulated annealing can occasionally accept a worse neighbouring solution to avoid becoming trapped too early, while tabu search uses memory to discourage cycling through recently visited moves. Their combination is designed to balance exploration and exploitation. Hybridization may improve search behaviour, but it also introduces design decisions about neighbourhoods, stopping rules and parameter settings. (Carvalho et al., 2026) The study evaluates the hybrid on a 40-stock S&P 500 subset, uses 2020–2023 for training and 2024 for testing, and compares against other metaheuristics and exact mixed-integer benchmarks; these design choices are more informative than a single in-sample best result (Carvalho et al., 2026).

Gad's systematic review of particle swarm optimization (PSO) provides broader context for this family of methods. PSO updates candidate solutions using information about individual and group search performance. Its appeal includes a relatively intuitive search mechanism and adaptability to different objective functions. However, the systematic review is not itself a portfolio-performance experiment; it should be used as background on the algorithm rather than as direct evidence that PSO outperforms portfolio-specific alternatives. (Gad, 2022)

Khan, Mohammed and Li propose a neurodynamic approach based on spiking neural networks for portfolio optimization. Their paper addresses computationally demanding formulations involving practical restrictions such as transaction costs and cardinality limits. This represents a different direction from conventional population-based heuristics: instead of relying solely on evolutionary operators or neighbourhood moves, the method uses a neural-dynamic system to seek a solution. It expands the available computational approaches, but implementation complexity, parameter interpretation and comparison against established solvers remain relevant considerations. (Khan et al., 2025)

4.5 Monte Carlo simulation and integrated approaches

Noaman, El-dosuky and Karawia combine Monte Carlo simulation with Operations Research for financial portfolio optimization. Monte Carlo methods generate many possible outcomes from specified assumptions and can be used to explore uncertainty in portfolio returns. Simulation is particularly helpful when a closed-form result is unavailable or when an analyst wants to inspect a distribution of possible outcomes rather than one expected value. (Noaman et al., 2020)

Simulation alone, however, is not necessarily an optimizer. It needs a mechanism for generating or selecting portfolio weights, an objective for comparing candidates and a procedure for checking constraints. Results are also affected by the assumed return distributions, the number of simulated scenarios and random variation. A strong workflow therefore uses simulation to evaluate uncertainty while retaining an explicit optimization model for the allocation decision.

V. Comparative Analysis and Discussion

5.1 From financial objectives to OR models

The modelling process starts by defining the decision. Let x_i represent the fraction of capital invested in asset i . A full-investment constraint can be written as the sum of all x_i equalling one. Long-only restrictions require non-negative weights, while upper bounds can limit exposure to any one asset. Expected portfolio return is typically a weighted sum of asset expected returns. Risk can be represented by variance, tail-loss measures, utility-based constraints, risk contributions or an efficiency score.

The selected objective then determines the optimization family. A covariance-based variance objective produces a quadratic term. A scenario-based CVaR objective can often be linearized using auxiliary variables. A maximum number of holdings requires binary selection variables and links those variables to the portfolio weights, creating a mixed-integer model. Robust optimization adds constraints or objectives that account for a set of plausible parameter values, while multistage stochastic programming represents decisions and uncertainty over time.

This matters because an algorithm should not be chosen independently of the model. A solver that is ideal for a convex QP may not directly handle a non-convex mixed-integer formulation. Conversely, a metaheuristic can be adapted to discrete constraints but may return a good feasible solution without a proof that no better solution exists. Good practice is to select the simplest model that represents the decision adequately and then use a solution method suited to its mathematical structure.

5.2 Comparative summary of the main approaches

Approach	What it emphasizes	Main strength	Main limitation
Mean–variance / QP	Expected return and total variance	Clear formulation; efficient-frontier interpretation	Sensitive to estimates; treats upside and downside deviations symmetrically
VaR / MIQP	Loss threshold at a chosen confidence level	Aligns with threshold-based risk limits	Optimization may be computationally difficult; VaR alone says little about tail severity
CVaR / LP	Average or expected loss in the tail	Scenario-based formulations can be tractable and interpretable	Depends on scenarios, confidence level and tail-loss specification
Robust optimization	Performance across plausible inputs	Explicitly accounts for estimation uncertainty	Uncertainty set calibration can make results too conservative or too optimistic

Multistage stochastic	Sequential decisions as information arrives	Represents long-term decisions and policy/scenario uncertainty	Large scenario trees and modelling burden
Fuzzy / possibilistic / neutrosophic	Imprecise or ambiguous inputs	Offers alternatives when point estimates or probability models are inadequate	Parameter interpretation and validation can be difficult
DEA with fuzzy returns	Relative efficiency of portfolios	Evaluates portfolios across multiple measures	Depends on comparison set and selected inputs/outputs
Metaheuristics	Difficult discrete or non-convex search spaces	Flexible for cardinality and complex constraints	No general global-optimality guarantee; tuning and runtime matter
Monte Carlo + OR	Distribution of outcomes under assumptions	Makes uncertainty visible through scenarios	Simulation quality depends on assumptions and is not an optimizer by itself
Neurodynamic methods	Computational solution of constrained models	Potential alternative for complex formulations	Implementation and comparative validation are important

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

5.3 How the risk measure changes the portfolio

The choice of risk measure is not a technical detail that can be separated from investment preferences. Variance captures dispersion around the mean and rewards diversification through covariance. CVaR places more emphasis on the losses that occur in the tail. VaR asks whether a specified loss threshold is likely to be breached. Utility-based risk measures incorporate preferences through a utility function. Risk parity allocates according to risk contributions rather than simply maximizing expected return for a given variance.

Consequently, two optimization methods can produce different portfolios even when they use the same asset data. A portfolio optimized for variance may include assets that diversify ordinary fluctuations but still experience large losses in stress scenarios. A CVaR model may allocate differently if a small number of severe outcomes dominate the tail. A robust model may reduce exposure to assets whose expected returns look attractive but are estimated imprecisely. The appropriate question is therefore not only ‘Which algorithm is best?’ but also ‘Which risk concept best represents the investor’s objective?’

5.4 Exact optimization versus heuristic search

Exact mathematical programming is attractive when the model can be formulated in a tractable structure. It offers reproducibility and, where the relevant conditions hold, a certificate of optimality. LP and convex QP solvers can be highly effective for large problems when the objective and constraints fit those structures. The PortfolioOptim package and the CVaR formulation developed by Rockafellar and Uryasev illustrate how mathematical reformulation can make risk-sensitive portfolio models computationally manageable.

Heuristic search becomes more attractive when discrete constraints, non-convexities or very large dimensions make exact methods difficult. Genetic algorithms, PSO, bee-colony search, biogeography-based optimization and simulated annealing–tabu search can explore complex feasible regions. Their flexibility is valuable, but they should be evaluated using multiple runs, transparent stopping rules, feasible-solution checks and comparable computing budgets. A single strong run does not establish consistent superiority.

A fair comparison should also distinguish solution quality from model quality. If a model uses poor return estimates or omits transaction costs, a more powerful algorithm can optimize the wrong problem more effectively. Similarly, a heuristic should not be criticized for failing to find a global optimum if the comparison method also lacks a known optimum. Where possible, small instances can be solved exactly to provide a benchmark for heuristic performance.

5.5 Connection with Operations Research coursework

The papers map directly onto several core OR concepts. Linear programming appears in scenario-based downside-risk formulations and other models with linear objectives and constraints. Quadratic programming is central to covariance-based mean–variance optimization. Mixed-integer programming is needed when decisions include whether an asset is selected, whether a minimum holding is opened, or whether transaction-related decisions are discrete. Stochastic programming models uncertainty across scenarios or time stages.

Monte Carlo simulation links optimization with probabilistic modelling and sensitivity analysis. Metaheuristics illustrate heuristic search, local improvement and the exploration–exploitation trade-off. Robust optimization shows how uncertainty can be incorporated into constraints and objectives. Across all these methods, sensitivity analysis is essential: changing expected returns, covariance estimates, risk-aversion parameters, confidence levels, uncertainty sets or algorithm settings may change the allocation.

This connection also provides a practical decision sequence for an analyst: define the investor’s objective; identify the relevant risk measure; specify feasible constraints; choose a mathematical formulation; select a suitable solver; validate feasibility and stability; and evaluate performance out of sample. The optimization result is a decision aid, not a guarantee of future returns.

VI. Results and Discussion

The comparison produces several broad findings. First, the risk measure and the solver must be considered together. CVaR is attractive partly because it can be formulated in ways that support linear optimization, while VaR can lead to a harder mixed-integer quadratic formulation. Mean–variance remains a useful benchmark, but it does not capture every investor’s concern about losses, parameter uncertainty or portfolio concentration. (Rockafellar & Uryasev, 2002; Cesarone et al., 2023; Palczewski, 2018)

Second, constraints change the nature of the problem. A portfolio with continuous weights and a budget constraint may be relatively tractable, whereas a portfolio with a maximum number of holdings, minimum investment thresholds or transaction costs may become a mixed-integer or non-convex problem. The studies on bee-colony algorithms, hybrid simulated annealing–tabu search, evolutionary risk parity, biogeography-based optimization and neurodynamic methods all respond to this broader class of computational challenges. (Mansourinia & Momeni, n.d.; Carvalho et al., 2026; Wijaya & Sidarto, 2023)

Third, there is no single representation of uncertainty that is suitable for every context. Robust optimization is appropriate when plausible parameter ranges can be defended. Multistage stochastic models are relevant when decisions unfold over time and later actions can respond to new information. Fuzzy, possibilistic and

neutrosophic models offer ways to express imprecision, but they require clear parameter interpretation. Monte Carlo simulation can illustrate the consequences of assumptions but cannot compensate for unrealistic scenario generation. (Georgantas et al., 2021; Gatzert et al., 2021; Bhattacharyya, 2013; Khalifa et al., 2022)

Fourth, computational performance should not be confused with investment performance. An algorithm may find a feasible portfolio quickly, but that alone does not show that the portfolio will perform well out of sample. Conversely, a model that performs well in a historical sample may be overfitted to that period. A balanced evaluation should consider expected return, volatility, downside risk, drawdown, concentration, turnover, transaction costs, runtime and stability across estimation windows.

Finally, cross-paper comparisons remain limited by differences in datasets and experimental design. The reviewed papers do not provide a single controlled tournament in which every method is evaluated using identical assets, dates, constraints and metrics. The evidence therefore supports a conditional conclusion: different methods are useful for different modelling requirements, but a universal performance ranking would go beyond what the source set establishes. (Georgantas et al., 2021; Liu & Xiao, 2021)

VII. Research Gaps and Future Scope

The most important gap is the lack of a common comparison protocol. Studies use different markets, periods, assets, benchmarks and constraints, so their reported outcomes cannot simply be placed side by side as if they were observations from one experiment. Future work should separate evidence about a model's mathematical properties from evidence about its empirical performance.

A practical extension would compare three representative models on the same asset universe: a mean–variance quadratic program, a CVaR-based linear program and a cardinality-constrained model solved with a metaheuristic. The experiment should use rolling estimation windows and an out-of-sample evaluation period. Each model should face the same budget, weight bounds and rebalancing assumptions. If the cardinality-constrained model is included, the maximum number of holdings and minimum weights should be identical across all relevant algorithms.

The evaluation should report annualized return, volatility, Sharpe ratio, maximum drawdown, CVaR at a specified confidence level, turnover, transaction costs and computation time. Results should be tested under different estimation windows and risk-aversion parameters. For stochastic algorithms, repeated runs should be reported rather than only the best run. An equal-weight portfolio and a simple benchmark should be included to show whether added model complexity provides meaningful improvement.

A further extension could test robust optimization or fuzzy-return models when the data and uncertainty assumptions are available. Such an extension should explain how uncertainty sets or membership functions are chosen and should test whether conclusions remain stable under reasonable alternatives. The goal would not be to make the model as complicated as possible, but to determine whether additional complexity improves the quality and reliability of decisions.

VIII. Conclusion

The selected literature shows that Operations Research offers a broad toolkit for portfolio optimization, from classical mean–variance programming to downside-risk models, robust and stochastic formulations, fuzzy approaches and modern heuristic or neurodynamic solvers. Each family addresses a different aspect of the problem. Mean–variance is clear and useful as a benchmark; CVaR focuses on tail losses and can often be optimized through LP; VaR can reflect threshold-based risk requirements but may

increase computational difficulty; robust and multistage stochastic models represent uncertainty more explicitly; and metaheuristics are flexible when discrete constraints make exact optimization challenging.

The review does not identify a universally superior method, because the studies do not share a common empirical design and their objectives differ. Instead, the strongest practical lesson is to match the model to the decision problem. The investor's risk preference, the quality of available data, the structure of constraints and the need for a provable optimum should guide the choice of formulation and solver.

A rigorous future comparison should use the same assets, sample windows, constraints, benchmarks and performance measures for every method, with out-of-sample testing and sensitivity analysis. That design would help distinguish genuine improvements from differences caused by data or assumptions. In this sense, Operations Research is most valuable not as a search for one universally best algorithm, but as a disciplined way to make portfolio decisions explicit, testable and better aligned with the risks investors actually face.

References

- [1] Bhattacharyya, R. (2013). Possibilistic Sharpe ratio based novice portfolio selection models. In R. Bhattacharyya et al. (Eds.), *ACER 2013* (pp. 33–45). CS & IT-CSCP. <https://doi.org/10.5121/csit.2013.3204>
- [2] de Simões Carvalho, S. M., Peixoto, M. da Silva, & de Freitas Jr., C. M. (2026). Simulated annealing–Tabu search integration for downside-risk portfolio optimization with cardinality constraints. *Journal of Heuristics*, 32, Article 27. <https://doi.org/10.1007/s10732-026-09601-9>
- [3] Cesarone, F., Martino, M. L., & Tardella, F. (2023). Mean-Variance-VaR portfolios: MIQP formulation and performance analysis. *OR Spectrum*, 45, 1043–1069. <https://doi.org/10.1007/s00291-023-00719-x>
- [4] Gad, A. G. (2022). Particle swarm optimization algorithm and its applications: A systematic review. *Archives of Computational Methods in Engineering*, 29, 2531–2561. <https://doi.org/10.1007/s11831-021-09694-4>
- [5] Gatzert, N., Martin, A., Schmidt, M., Seith, B., & Vogl, N. (2021). Portfolio optimization with irreversible long-term investments in renewable energy under policy risk: A mixed-integer multistage stochastic model and a moving-horizon approach. *European Journal of Operational Research*, 290(2), 734–748. <https://doi.org/10.1016/j.ejor.2020.08.033>
- [6] Geissel, S., Graf, H., Herbinger, J., & Seifried, F. T. (2022). Portfolio optimization with optimal expected utility risk measures. *Annals of Operations Research*, 309, 59–77. <https://doi.org/10.1007/s10479-021-04403-7>
- [7] Georgantas, A., Doumpos, M., & Zopounidis, C. (2021). Robust optimization approaches for portfolio selection: A comparative analysis. *Annals of Operations Research*. <https://doi.org/10.1007/s10479-021-04177-y>
- [8] Hochreiter, R. (2015). An evolutionary optimization approach to risk parity portfolio selection. arXiv preprint arXiv:1411.7494.
- [9] Khalifa, H. A. E.-W., Saeed, M., Rahman, A. U., & El-Morsy, S. (2022). An application of pentagonal neutrosophic linear programming for stock portfolio optimization. *Neutrosophic Sets and Systems*, 51, 653–665. <https://doi.org/10.5281/zenodo.7135397>

- [10] Khan, A. H., Mohammed, A. M., & Li, S. (2025). Portfolio optimization: A neurodynamic approach based on spiking neural networks. *Biomimetics*, 10(12), Article 808. <https://doi.org/10.3390/biomimetics10120808>
- [11] Liu, S., & Xiao, C. (2021). Application and comparative study of optimization algorithms in financial investment portfolio problems. *Mobile Information Systems*, 2021, Article 3462715. <https://doi.org/10.1155/2021/3462715>
- [12] Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91. <https://doi.org/10.2307/2975974>
- [13] Mansourinia, M., & Momeni, A. (n.d.). Cardinality constrained portfolio optimization using bee colony algorithm (case study: Tehran Stock Exchange). [Publication details not stated in the supplied PDF.]
- [14] Noaman, N. M., El-dosuky, M. A., & Karawia, A. (2020). Financial portfolio optimization using Monte Carlo and Operation Research. *International Journal of Computer Applications*, 175(34), 43–46.
- [15] Palczewski, A. (2018). LP algorithms for portfolio optimization: The PortfolioOptim package. *The R Journal*, 10(1), 308–327.
- [16] Rockafellar, R. T., & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7), 1443–1471.
- [17] Wijaya, W., & Sidarto, K. A. (2023). Large-scale portfolio optimization using biogeography-based optimization. *International Journal of Financial Studies*, 11(4), Article 125. <https://doi.org/10.3390/ijfs11040125>
- [18] Zhou, Z., Chen, E., Xiao, H., Ren, T., & Jin, Q. (2019). Performance evaluation of portfolios with fuzzy returns. *RAIRO—Operations Research*, 53, 1581–1600. <https://doi.org/10.1051/ro/2019071>