Proceedings of the 7th International Conference on Economics and Social Sciences (2024), ISSN 2704-6524, pp. 682-691

The 7th International Conference on Economics and Social Sciences
Exploring Global Perspectives:
The Future of Economics and Social Sciences
June 13-14, 2024
Bucharest University of Economic Studies, Romania

Enhancing Portfolio Structure with Evolutionary Multi-Objective Optimisation

Robert-Ștefan CONSTANTIN [1], Marina-Diana AGAFIȚEI [2],
Adriana AnaMaria DAVIDESCU[3*]

DOI: 10.24818/ICESS/2024/061

Abstract

In this study, we define the criteria for fund allocation in an investment portfolio based on three key issues: maximizing returns, minimising risk, and optimal asset allocation. The context of solving these issues reveals that the best solutions are not those that sequentially maximise or minimise each criterion but rather those that achieve an optimal compromise between them, known in the specialised literature as the Pareto front.

To identify a set of nondominated solutions, we utilise a specialized evolutionary algorithm for multi-objective optimisation, the Nondominated Sorting Genetic Algorithm II (NSGA-II). This is a fast and elitist evolutionary algorithm based on a process of sorting and selecting the best agents for the repopulation of new solving sets. By using this algorithm, we generate different sets of possible solutions, also testing various mutation rates of the agents to study different approaches to favourable combinations for fund allocation. The subjects of these iterations will be a set of some of the most successful assets listed on the Bucharest Stock Exchange, simultaneously including a considerable part of the Bucharest Exchange Trading Index, over a period that encompasses both the COVID-19 pandemic and the Ukrainian war shocks. Subsequently, we evaluate the performance of these portfolio weights over time, analysing their performance and identifying differences in the evolutionary genome behaviour in comparison to the traditional Markovitz method of quadratic mean-variance equation.

Keywords: Evolutionary Multi-Objective Algorithm, NSGA-II, Portfolio, Risk, MOEA, MOOP.

JEL Classification: G11, C61, G17, G12, C63.

[1] Bucharest University of Economic Studies, Bucharest, Romania, constantinrobert21@stud.ase.ro.
[2] Bucharest University of Economic Studies, Bucharest, Romania, diana.agafitei@csie.ase.ro.
[3] Bucharest University of Economic Studies, Bucharest, Romania, adriana.alexandru@csie.ase.ro, National Scientific Research Institute for Labour and Social Protection, adriana.davidescu@incsmps.ro.
* Corresponding author.

1. Introduction

When speaking of financial time series, a set of particular characteristics must be taken into consideration, precisely their stochastic nature, given by a high sensitivity to shocks, the existence of volatility clustering, nonstationary structure, and ultimately inconsistent variability over time. Having said this, we can conclude the unpredictable nature of financial markets, ranging from sudden changes in prices and returns to unexpected changes due to policy decisions or market sentiments.

To address such uncertainty issues, analysts and investors use vast quantitative models and statistical frameworks to continuously validate market models and identify vivid relations between different assets and the market.

One way of preparing financial data for modelling is working with returns, discrete or continuous, instead of stock prices. Thus, we obtain a more desirable way to understand gains and an easier set of statistical properties to work with, as Campbell and MacKinlay (1997) highlight. The advantages of working with returns can be associated with operating within the risk dimension, obtaining the ability to directly observe the level of volatility in a given period while maintaining a data set with an average value and standard deviation close to 0 and 1.

Past research and theorems assumed that financial markets are a stable economic environment and that all investors have an equal level of utility and risk aversion. In fact, this has been shown on many occasions to be otherwise. Independent of the specific problem at hand, one general problem which has been given much attention concerns how to optimally estimate portfolio weights. In particular, the problem in portfolio risk optimization consists of competing objectives that must be optimized, involving maximizing returns and minimizing risk. These two main objectives have been mathematically formalized by the mean-variance optimization problem of Markowitz in 1960 and have since then been extended to include a variety of constraints and additional considerations. Once again, the inadequacy of this method lies in the optimization under a set of traditional market hypotheses.

The traditional optimization methods very often fail when we add real-life constraints. Already limiting the number of assets to invest in, which we will be considering in our further analysis, creates major problems for these methods. They might even turn computationally infeasible with the rising complexity of the constraints, turning even some of the cases into NP-hard (Weilong, 2023).

With the enhancements made in computational processes, we are getting closer and closer to finding better ways of exploring the interactions between assets and financial markets. As noted above, our approach lies within the realm of evolutionary algorithms, which turned out to be very effective in solving computationally intensive tasks. Of these, the Nondominated Sorting Genetic Algorithm II is chosen in this work due to its remarkable efficiency in handling multi-objective optimisation problems with considerable efficiency. (Deb et al, 2002). NSGA-II utilises a rapid non-dominated sorting method and a crowding distance mechanism to preserve solution diversity, making it well-suited for financial portfolio optimisation.

As stated above, the constraints of our objectives are taken from the biobjective problem of mean-variation of Markovitz and, lastly, the taking into consideration of cardinal assets. Thus, we obtain a tri-objective optimisation issue. Numerous studies also took an approach to this matter by using evolutionary algorithms, noteworthy for this article that as inspiration the research of Anagnostopoulos & Mamanis (2010), in which they observe the performances of 3 MOEA's methods (Multi Objective Evolutionary Algorithms). In our case, we will further observe how NSGA-II will compare with one traditional technique based on the quadratic mean-variation estimation of assets. By examining both approaches, we aim to highlight the advantages and limitations of the traditional efficient frontier method and demonstrate how evolutionary algorithms can offer robust solutions in the realm of financial portfolio optimisation.

In the following sections, we will provide a brief explanation of how the constraints are integrated into the objective function, the operational mechanisms of genetic algorithms within this framework, and the computational methodology of our comparison counterpart.

2. Multi-Objective Portfolio Optimisation

In portfolio theory, the concept of an efficient portfolio is central to optimising returns while managing risk. According to the seminal work of Huang & Litzenberger (1988) an efficient portfolio is defined as one that lies on the portfolio frontier with an expected rate of return strictly higher than that of the minimum variance portfolio. To achieve that, we make use of the following constraints:

$$\min \rho\,(\mathbf{x}) = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij} \tag{1}$$

$$\max \mu\,(\mathbf{x}) = \sum_{i=1}^{n} x_i \mu_i \tag{2}$$

For the formulations above, σ represents the variation and covariation of our assets, respectively, the risk they contain and μ symbolizes the expected return of one asset. Furthermore, the efficient portfolio is also limited by the fact that one must expend all his funds into this set of securities, therefore $\sum_i x_i = 1$.

To further refine the portfolio, we incorporate the cardinality constraint aimed at minimising the number of assets with nonzero weights, thereby simplifying the portfolio and forcing the algorithm to choose the best ratio of risk and return for an assett.

$$min\ card(\mathbf{x}) = \sum_{i=1}^{n} 1_{x_i} > 0 \tag{3}$$

Evolutionary algorithms (EAs) are methods inspired by real-life processes, specifically the idea of natural selection and evolution (Coello, 2007). MOEAs are characterised by generating a set of Pareto-front solutions, focussing on finding the best non-dominated outcomes. This approach ensures that all constraints are partially satisfied without having any single solution dominated by another. With this diversity we also obtain different ways the algorithm can build a portfolio, seeing how for every iteration a different approach is taken.

An MOEA is called an elitist algorithm when it is able to generate new populations based on combining the best performing agents of past iterations. By doing this, each generation is slightly improving; later, we find an optimum point or we limit the number of simulations.

3. Quadratic Problem for Mean-Variance

The mean-variance optimisation framework is widely used in finance to construct efficient portfolios that provide the best possible return for a given level of risk (Elton, 2014). It can be formulated as a quadratic programming problem with the objective of minimising the portfolio variance subject to constraints on the expected return and the portfolio weights. The problem is states as the following (Huang & Litzenberger, 1988):

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^{\top} \Sigma \mathbf{w} \tag{4}$$

$$\mathbf{w}^{\top} \mu = p \tag{5}$$

$$\mathbf{w}^{\top \mathbf{1}} = 1 \tag{6}$$

It is stated that a portfolio is of frontier if and only if the weight vector of the portfolio is the solution of the quadratic equation. Consequently, we have the constraint for minimising risk, where w is the, then the p being the targeted expected return of portfolio.

$$\mathcal{L}(\mathbf{w}, \lambda_1, \lambda_2) = \frac{1}{2} \mathbf{w}^{\top} \Sigma \mathbf{w} - \lambda_1 (\mathbf{w}^{\top} \mu - p) - \lambda_2 (\mathbf{w}^{\top \mathbf{1}} - 1) \tag{7}$$

The Lagrangian function incorporates the objective function and the constraints using Lagrange multipliers (λ_1 and λ_2). This formulation helps in solving the optimisation problem by finding the stationary points of the Lagrangian, which satisfy both the objective function and the constraints.

Although the mean-variance optimisation formula pioneered in Modern Portfolio Theory, a set of significant drawbacks had been associated with it (Haugh, 2016). It is highly sensitive to input estimates of expected returns and covariances, at the same time being incapable of taking into account asymmetry and kurtosis, making the optimal asset allocations prone to significant changes with small estimation errors. Additionally, this is a static model that only relies on historical data for estimating only one outcome for asset allocation. Compared to our testing counterpart has a dynamic and can estimate different scenarios using the same initial dataset.

4. Empirical Results

Our database consists of six assets, some of the best performing ones included in the Bucharest Exchange Trade Index (BET), having a proportion of around 60% of it. They have the following symbols: BRD, TLV, SNP, SNG, SNN, EL. BRD and TLV belong to the banking sector, while the other classify in the energy sector.

For this study, we used a period of approximately 5 years, from 2019/01/04 to 2024/03/19. This time axis also includes the effects of the COVID-19 pandemic and the Ukrainian war, thus having an even more volatile set of dates.

It is also worth mentioning that this dataset was obtained from the Bucharest Stock Exchange (BVB) through a formal request for academic use of the information, ensuring the accuracy and relevance of the information used in this study.

Figure 1. Daily Returns of Assets

Source: Self-computed in RStudio using Plotly library (2024).

Before computing the algorithm, we took a step back to analyse the average return of each series and its level of volatility. By visualising the monthly returns using boxplots, we observed that the mean returns tend to be positive, even though there were periods of increased volatility, particularly around significant global events such as the COVID-19 pandemic and the Ukraine war. This increased volatility can be attributed to the phenomenon of volatility clustering, as noted by Mandelbrot (1963), where periods of high uncertainty tend to be followed by even higher uncertainty. Letting the algorithm experience these effects will also demonstrate its ability to avoid risk and capacity to evaluate assets.

Figure 2. Monthly Returns of Assets. Boxplots representation

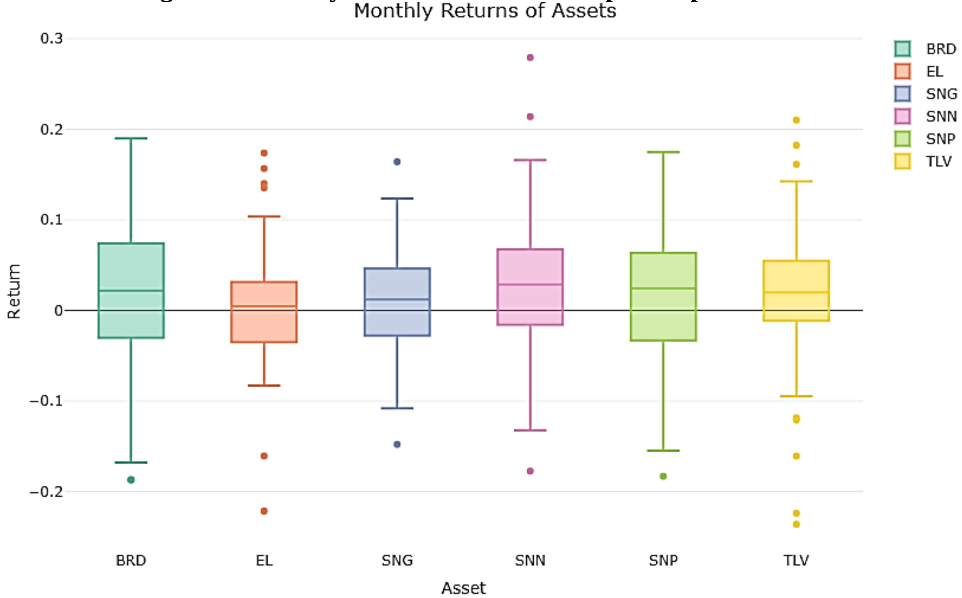

Source: Self-computed in RStudio using Plotly library (2024).

For the NSGA-II method, we performed three sets of simulations across four different mutation rate settings. In each simulation, a total of 1000 iterations were set, with a standard population size of 100. The crossover probability was maintained at 0.9 in all scenarios and mutation rates were established at 0.01, 0.05, 0.10, and 0.15.

Computational analysis was performed using RStudio, using the “nsga2R” library (Tsou, 2022) to execute NSGA-II simulations and the ‘series’ package (Trapletti, 2024) for traditional estimates. This implementation was inspired by the methodology outlined in Adyatama's work published in RPubs (2021), with the distinction of excluding a risk-free rate. The risk-free rate would typically act as a baseline for expected returns on risky assets, as they would be derived from governmental bonds (Damodaran, 2008). As a result, exclusion was necessitated by the extended time period over which the model was tested.

Table 1. Estimated return of simulated portfolios

Total returns	0.01	0.05	0.1	0.15
I	0.00037	0.00064	0.00068	0.00065
II	0.00088	0.00091	0.00091	0.00067
III	0.00092	0.00054	0.00066	0.00063

Source: Self-computed in RStudio, 2024.

Table 2. Estimated risk of simulated portfolios

Risk	0.01	0.05	0.1	0.15
I	0.00058	0.00014	0.00016	0.00015
II	0.00014	0.00017	0.00032	0.00065
III	0.00014	0.00018	0.00014	0.00016

Source: Self-computed in RStudio, 2024.

In the first table, we have arranged the results of simulated portfolio returns based on their number of simulations and the rate on mutation. This was done to study the randomness of each computation. In the following table we have the risk correspondence of each simulation. From this set we obtained both satisfying combinations of return/risk and lesser performing variants. Based on this, we have observed some patterns that suggest that the algorithm undergoes the most significant changes during the initial iterations of the simulated environment. Such behaviour can be attributed to the algorithm's exploration phase. As the iterations progress, the algorithm tends to exploit the most promising regions of the solution space, leading to a reduction in variability and a gradual convergence towards optimal or near-optimal solutions.

In addition, we decided to extract the best options from each mutation rate value based solely on the performance of the returns. We did this because we want to study why the algorithm took such paths even when confronted with significant levels of risk, like the case of simulation II.0.15 where the risk is almost the same as the return, respectively, 0.00067 and 0.00065.

From the results shown in Table 3, it appears that with a higher mutation rate, without being restricted in terms of minimum and maximum fund allocation for one asset, the algorithm tends to distribute most of its resources into two securities. Specifically, as the mutation rate increases from 0.01 to 0.15, we observe a significant concentration of weights in the SNN and SNP assets.

This phenomenon can be attributed to the exploration-exploitation trade-off inherent in evolutionary algorithms. At higher mutation rates, the algorithm introduces more substantial changes in each iteration, increasing the diversity of the portfolio compositions. However, this can also lead to a higher likelihood of converging on a few highly dominant solutions, especially if those solutions appear to offer superior return-risk trade-offs early on.

Table 3. Weights of portfolios

	0.01	0.05	0.10	0.15	M-V equation
BRD	0.037	0.024	0.175	0.026	0.087
EL	0.134	0.009	0.059	0.223	0.200
SNG	0.102	0.030	0.107	0.024	0.261
SNN	0.315	0.235	0.521	0.662	0.188
SNP	0.386	0.238	0.104	0.049	0.177
TLV	0.026	0.464	0.034	0.016	0.088

Source: Self-computed in RStudio, 2024.

In our case, the asset SNN seems to dominate the portfolio allocations at higher mutation rates. This could be due to these assets demonstrating favourable return characteristics that align well with the algorithm's objectives. The mutation rate of 0.15, in particular, shows a very pronounced allocation to SNN with a weight of 66.2%.

On the other hand, the mean-variation method tends to spread the weights more evenly across multiple assets, with no single asset substantially dominating the portfolio.

Figure 3. Wealth index representation

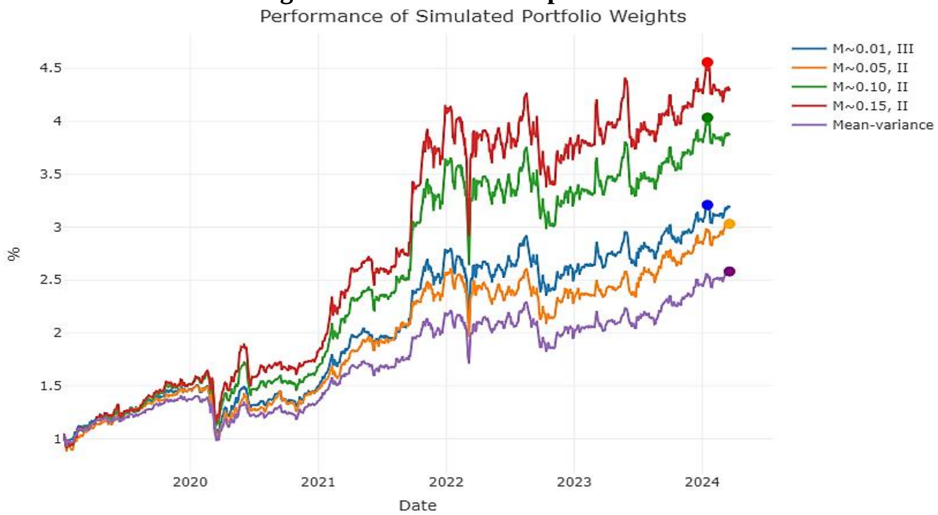

Source: Self-computed in RStudio using Plotly library (2024).

Lastly, we have generated five different portfolios based on the weight we obtained through simulation and the mean-variance method and tested their performance on the time axis. The wealth index was chosen to reflect how many times the portfolio grow in a time span while testing different structures. The y axis represents the percentage growth rate (x times, 1.5 = 150%). It is also noted that the output of the mean variance formula was a return of portfolio of 0.00075 and an overall risk of 0.00012.

On a general note, NSGA-II managed to select more rewarding portfolios, even if all the simulations had a bigger risk ratio, some of them with a small increase, the ones with a mutation rate of 0.01% and 0.05%, they still found a better allocation of funds for a greater income. This shows how the mean-variance is a lot more focused on minimal risk rather than truly identifying the best option in fund allocation.

It is also notable that the portfolios had the best overall performance, earning a multiplication of 4.56 times of initial investment for the highest mutation rate (red line) and the second highest of 4.03 (green line). Even if the return value of the initial simulation was 0.00092 for the green line and 0.00067 for the red, those results were caused by the higher risk rates of 0.00032 and 0.00065.

5. Conclusions

The NSGA-II algorithm turned out to be a great solution for multiobjective optimization problems, thus creating different efficient scenarios with time and even handling higher ratios of risk. It arrives at high-return options and at the same time also gives some safer and lower-risk alternatives, therefore attracting many

Investment strategies.

In the present research, all the sets were created using a formula similar to the mean-variance equation with nonzero asset allocations. This flexibility makes NSGA-II search for an even larger space of feasible solutions and hence proves to be better in portfolio optimization.

We amassed enough evidence to suggest that multi-objective evolutionary algorithms can detect even more complex solutions, given further research on constraints and adaptability. On the other hand, sensitivity to parameter settings and tendencies towards premature convergence are defects that cannot be ignored in NSGA-II. These limitations suggest that perhaps our model was not fitted for this portfolio performance test.

Integration of long-term and short-term memory methods in stock behavior, dynamic portfolio structure over time, and conditional volatility would improve the performance of algorithms. Integrating methods for identifying long-term and short-term memory in stock behaviour, dynamically altering portfolio structure over time, and taking into account conditional volatility could enhance algorithm performance. One way would be the integration of more advanced Machine Learning (ML) techniques like ARIMA or SARIMA for testing and identification of recurrent patterns or seasonal windows or even augmentation via neural networks (NN). These advancements could lead to more precise and reliable portfolio optimisation, benefiting investors with better-informed strategies.

Acknowledgement

The research study was elaborated within the Data Science Research Lab for Business and Economics of the Bucharest University of Economic Studies, within the project "Analysis of the economic recovery and resilience process in Romania in the context of sustainable development," contract no. 1353/21.06.2024 and the project "The development and adaptation of collaborative processes in the excellent research conducted at ASE Bucharest, in the context of modern challenges brought by Open Science and Artificial Intelligence (eXROS)," contract: CNFIS-FDI-2024-F-0302/27.05.2024.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author(s) utilised ChatGPT to enhance readability and language clarity. Following the use of this tool, the author(s) meticulously reviewed and revised the content as necessary and assumed (s) full responsibility for the final content of the publication.

Bibliography

[1] Adyatama, A. (2021). Multi-Objective Optimization with NSGA-II. Retrieved from Rpubs: https://rpubs.com/Argaadya/moea.

[2] Anagnostopoulos, K., Mamanis, G. (2010). A portfolio optimization model with three objectives and discrete variables. Computers & Operations Research, 37(7), 1285-1297. doi:10.1016/j.cor.2009.09.009.

[3] Campbell, J.Y. (1997). The Econometrics of Financial Markets. New Jersey: Princeton University Press.

[4] Coello, C.A. (2007). Evolutionary Algorithms for Solving Multi-Objective Problems. (E.D. Goldberg, J.R. Koza, Eds.) doi:10.1007/978-0-387-36797-2.

[5] Damodaran, A. (2008). What is the Riskfree Rate? A Search for the Basic Building Block. SSRN. doi:10.2139/ssrn.1317436.

[6] Deb, K., Pratap, A., Sameer, A.T.M. (2002). A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. IEEE Transaction on Evolutionary Computation, 6(2), 182-197.

[7] Elton, E.J. (2014). Modern Portfolio Theory and Investment Analysis. John Wiley & Sons.

[8] Haugh, M. (2016). Mean-Variance Optimization and the CAPM. Foundations of Financial Engineering.

[9] Huang, C., Litzenberger, R.H. (1988). Foundations for Financial Economics. NY: Elsevier.

[10] Mandelbrot, B.B. (1963). The variation of certain speculative prices. Journal of Business, 36, 392-417.

[11] Markowitz, H.M. (1960). Portfolio Selection: Efficient Diversification of Investments. The Journal of Finance, 15(3), 447. doi:10.2307/2326204.

[12] Trapletti, A. (2024). TimeSeriesAnalysisandComputationalFinance. Retrieved from https://cran.r-project.org/web/packages/tseries/tseries.pdf.

[13] Tsou, C.-S. (2022). Elitist Non-Dominated Sorting Genetic Algorithm. Retrieved from https://cran.r-project.org/web/packages/nsga2R/nsga2R.pdf.

[14] Weilong, L. (2023). Evolutionary multi-objective optimisation for large-scale portfolio selection with random and uncertain returns. Belfast: QMS Working Paper. doi:10.2139/ssrn.4376779.