

Applications of Operations Research in Investment Portfolio Allocation: A Literature-Based Analysis of Linear Programming Models

Operations Research Paper | Problem type: Linear Programming Problem (LPP)

Pragya Shukla¹, Ria Kaur Ahuja², Prisha Mantri³, Parth Chachan⁴, Rayhan Sheikh⁵



<https://doi.org/10.55041/ijssmt.v2i10.038>

Cite this Article: Shukla, P., Ahuja, R. K., Mantri, P., Chachan, P. & Sheikh, R. (2026). Applications of Operations Research in Investment Portfolio Allocation: A Literature-Based Analysis of Linear Programming Models. International Journal of Science, Strategic Management and Technology, 02(10), 1-9. <https://doi.org/10.55041/ijssmt.v2i10.038>

License:  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

Abstract

Investors with a fixed budget must decide how much to put into each of several stocks, balancing return against risk. The classical answer, Markowitz's mean-variance model, is a quadratic programme that becomes heavy as the number of stocks grows. This paper reviews two published studies that solve the same problem with linear programming instead. Papahristodoulou compares a quadratic model with a maximin model and a mean-absolute-deviation (MAD) model on five Stockholm stocks. Kucukbay and Araz build a portfolio from 89 Borsa Istanbul stocks using linear physical programming (LPP) and fuzzy goal programming (FGP). For each study we set out the real-life problem, the mathematical model, the solution method, the results and the weaknesses. In the first study the author reports the MAD portfolio as close to the quadratic one while avoiding the covariance matrix, although on the published figures its risk sits nearer to that of maximin. In the second, the LPP portfolio had the best Sharpe index in the test year (0.43, against 0.36 for FGP and 0.26 for the market index). Both studies rest on small samples and short test periods, so their results show that linear models are practical, not that they are best.

Keywords: operations research; linear programming; portfolio selection; mean absolute deviation; maximin; goal programming; Sharpe index

1. INTRODUCTION

Operations Research (OR) uses mathematical models to help people make good decisions when money, time, materials or people are limited. A typical study defines the problem, builds a model, solves it, checks that the answer is sensible, and then puts it to use. Linear programming (LP) is the most widely used of these tools. An LP has decision variables (the quantities we control), an objective function (what we want to maximise or minimise), constraints (the limits we must respect) and non-negativity conditions. It assumes that effects are proportional and additive, that quantities can be divided freely, and that the data are known with certainty. Such problems are solved with the simplex method or with software such as Excel Solver, LINGO or CPLEX.

Investment portfolio allocation fits this framework naturally. An investor has a fixed budget and a list of securities, and must decide how much to invest in each. Money is the scarce resource, so this is a resource allocation problem. Two quantities matter: return, the profit expected, and risk, the chance that the outcome is worse than expected. Spreading money over several stocks (diversification) lowers risk, because a poor performer can be offset by a better one. Table 1 lists the terms used in the rest of the paper.

Table 1. Key terms

Term	Meaning
Return	Percentage gain or loss on an investment over a period, for example one month.
Risk	How much returns vary; more variation means more risk.
Variance / standard deviation	The usual statistical risk measure. Squaring deviations makes a model quadratic.
Absolute deviation	Distance of a return from its average, ignoring sign. A risk measure that suits linear models.
Short selling	Selling shares one does not own. Both studies disallow it, so every investment amount is ≥ 0 .
Sharpe index	$(\text{Portfolio return} - \text{risk-free return}) \div \text{portfolio risk}$; return earned per unit of risk.
Out-of-sample test	Checking a model on new data not used to build it.

The aim of this paper is to examine how LP is used for portfolio allocation by analysing two published studies. The paper is organised as follows. Section 2 reviews the background literature and the two selected studies. Section 3 explains their mathematical models. Section 4 reports their results. Section 5 compares the studies and discusses their strengths and limits, and Section 6 concludes.

2. LITERATURE REVIEW

2.1 From variance to linear risk measures

Modern portfolio theory begins with Markowitz (1952). His point was that an investor should look at how a portfolio behaves as a whole, not at each stock alone, and he proposed minimising the portfolio variance for a required level of return. The idea has lasted, but it is awkward to use. Variance involves products of the form $x_i x_j$, so the objective is quadratic, and every pair of stocks needs a covariance estimate. For 300 securities that means roughly 45,000 variances and covariances (Papahristodoulou, 2002).

Konno and Yamazaki (1991) offered a way round this. They replaced variance with the mean absolute deviation of returns, which can be written as a linear programme, and tested the idea on the Tokyo stock market. Both studies we review follow this line of thinking: keep the spirit of Markowitz, but use models that ordinary LP software can solve quickly.

2.2 Papahristodoulou: a small, fully worked example

Papahristodoulou (2002) starts with a practical question. A portfolio manager has a client who will invest at most 100,000 SEK, wants at least 3% return per month, and will not put more than 75% of the money in one share. Which of five Stockholm shares should be bought, and in what amounts? The author argues that the quadratic model is slow for large problems, assumes normally distributed returns, and can even pick an inefficient portfolio. As an illustration, a portfolio earning 8% or 16% should always be preferred to one earning 7% or 8%, yet minimum variance prefers the second because it varies less. To offer an alternative, the paper builds two LP models, maximin and MAD, and compares them with a quadratic benchmark using twelve months of data.

The comparison is helpful because the numbers are small enough to check by hand. Maximin gave the highest return and the highest risk, the quadratic model the lowest of both, and the author describes MAD as sitting close to the quadratic model. The out-of-sample results are less tidy. Over the next six months all three portfolios lost money, so the test does not crown a winner. The author nevertheless concludes that "one can do much more with LP than many think

of". It is worth noting that a longer version of this work, written with Dotzauer and published in the *Journal of the Operational Research Society* (Papahristodoulou and Dotzauer, 2004), uses 67 securities over 48 months. There, the maximin portfolios turned out the most robust in the six-month test. That is not the pattern seen in the five-share example, which is a reminder that conclusions drawn from five stocks may not carry over to a full market.

2.3 Kucukbay and Araz: many goals, realistic rules

Kucukbay and Araz (2016) move from a classroom-sized example to a real market. They select from 89 Borsa Istanbul companies that traded continuously over 60 months (April 2009 to March 2014) and want both a high expected return and low risk. Because these goals pull in opposite directions, they use goal programming methods. Linear physical programming (LPP) lets the investor describe each goal in ranges, from "ideal" down to "unacceptable", instead of fixing one target. Fuzzy goal programming (FGP) gives each goal a satisfaction score between 0 and 1 and raises the lowest score as far as possible. The portfolios must also obey practical rules: no more than 25% in any sector, at least 50% in BIST 50 firms, and no more than 10% in a single stock.

The two methods were compared on a test year (April 2014 to March 2015) using the Sharpe index. LPP scored 0.43, FGP 0.36 and the BIST 100 index 0.26, so both portfolios beat the market and LPP did best. The authors describe LPP as a good alternative to FGP. A careful reader should keep in mind that this rests on one training window and one test year.

Table 2. The two selected studies

	Study 1	Study 2
Reference	Papahristodoulou (2002), working paper	Kucukbay and Araz (2016), <i>IJOCTA</i> 6(2), 121–128
Data	5 Stockholm stocks, 12 months	89 Borsa Istanbul stocks, 60 months
Methods	Quadratic programming, maximin LP, MAD LP	Linear physical programming, fuzzy goal programming
Validation	Next 3 and 6 months	Next 12 months, Sharpe index

2.4 What the literature leaves open

Taken together, the studies show that linear models can replace the quadratic one, but they leave gaps. Neither names enough about its software to be easily repeated, neither includes transaction costs, and both lean on short test periods. The second study also measures risk in a way that ignores how stocks move together, which we return to in Section 5. These gaps shape the discussion and the conclusion that follow.

3. PROBLEM FORMULATION AND MATHEMATICAL MODELS

This section sets out the models used in the two studies in the same notation, so that they can be compared.

3.1 Study 1: three models for five stocks

Let x_j be the amount (SEK) invested in share j ($j = A, \dots, E$), r_{jt} the return of share j in month t ($t = 1, \dots, 12$), and $\bar{r}_j$ its average return. The averages are 0.0207, 0.0316, 0.0323, 0.0337 and 0.0376 for shares A to E. The budget is 100,000 SEK, the minimum return is 3,000 SEK per month, and no share may exceed 75,000 SEK. Short selling and transaction costs are ignored.

Model A: Quadratic programming (Markowitz benchmark)

Here x_j is a fraction of the budget. The model is not linear and was solved with QSB+ software.

$$\text{Minimise } \sum_i \sum_j \sigma_{ij} x_i x_j \quad (1)$$

$$\text{subject to } \sum_j \bar{r}_j x_j \geq 0.03, \quad \sum_j x_j = 1, \quad 0 \leq x_j \leq 0.75 \quad (2)$$

Model B: Maximin LP

The idea is to make the worst month as good as possible. Let Z be the lowest monthly return of the portfolio.

$$\text{Maximise } Z \quad (3)$$

$$\text{subject to } \sum_j r_{jt} x_j - Z \geq 0 \quad (t = 1, \dots, 12), \quad \sum_j x_j \leq 100,000 \quad (4)$$

$$\sum_j \bar{r}_j x_j \geq 3,000, \quad 0 \leq x_j \leq 75,000 \quad (5)$$

Only the downside is controlled; upside variation is left free. A weakness is that if every share loses money in the same month, Z must be allowed to go negative or the model has no solution.

Model C: Mean absolute deviation (MAD) LP

Risk is the average absolute distance of the portfolio's monthly return from its average. Let $d_{jt} = r_{jt} - \bar{r}_j$. Because the absolute value is not linear, a helper variable $Y_t \geq 0$ is introduced for each month.

$$\text{Minimise } (1/12)(Y_1 + Y_2 + \dots + Y_{12}) \quad (6)$$

$$\text{subject to } Y_t + \sum_j d_{jt} x_j \geq 0, \quad Y_t - \sum_j d_{jt} x_j \geq 0 \quad (t = 1, \dots, 12) \quad (7)$$

$$\sum_j x_j = 100,000, \quad \sum_j \bar{r}_j x_j \geq 3,000, \quad 0 \leq x_j \leq 75,000 \quad (8)$$

Requiring Y_t to be at least both the deviation and its negative makes it settle at the larger of the two, which is exactly the absolute value, and the model stays linear. As an example, the January deviations of shares A to E are 0.0333, 0.0004, 0.0083, 0.0043 and 0.0114, which gives the pair $Y_1 \pm (0.0333A + 0.0004B + 0.0083C + 0.0043D + 0.0114E) \geq 0$. The model has $2 \times 12 + 2 = 26$ constraints. The number of constraints depends on the number of months, not on the number of shares, and an optimal solution holds at most $2T + 2$ shares.

To compare the three portfolios, the paper uses the utility $U = E(r) - w\sigma$, where w measures how much the investor dislikes risk ($w = 0$ means indifference to risk). It then follows the portfolios for the next 3 and 6 months using real returns.

3.2 Study 2: two goals, two methods

Let X_i be the proportion of the budget invested in stock i , r_i its expected monthly return and σ_i the standard deviation of its return, both estimated from 60 monthly observations for 89 companies. The two goals are:

$$\text{Goal 1 (maximise return): } Z_1 = \sum_i r_i X_i \quad (9)$$

$$\text{Goal 2 (minimise risk): } Z_2 = \sum_i \sigma_i X_i \quad (10)$$

Both methods share four constraints: (i) all money is invested, $\sum X_i = 1$; (ii) the stocks of any one sector total at most 0.25; (iii) at least 50% of the money goes to BIST 50 firms; and (iv) $0 \leq X_i \leq 0.10$.

Linear physical programming. The decision maker describes each goal using six ranges of satisfaction. Table 3 shows the target values used. Deviation variables measure how far a goal falls into a worse range, and the model minimises a weighted sum of these deviations. The weights are not chosen by hand: an algorithm generates them so that worse ranges carry larger weights, which pushes the solution strongly away from unacceptable outcomes.

Table 3. Target ranges for LPP (Kucukbay and Araz, 2016, Table 2)

Satisfaction level	Expected return (%)	Risk (%)
Ideal	> 4.059	< 10.340
Desirable	4.059 – 3.247	10.340 – 12.408
Tolerable	3.247 – 2.760	12.408 – 13.648
Undesirable	2.760 – 2.354	13.648 – 14.688
Highly undesirable	2.354 – 2.140	14.680 – 16.320
Unacceptable	< 2.14	> 16.320

The original table gives 14.688 and 14.680 at the boundary between two ranges; we reproduce it as printed.

Fuzzy goal programming. Each goal gets a membership function μ between 0 (not satisfied) and 1 (fully satisfied). For return, μ rises from 0 at 2.148 to 1 at 4.059; for risk it falls from 1 at 10.347 to 0 at 16.326. These limits come from the pay-off table (Table 4), in which each goal is first optimised alone. The model then maximises the lowest satisfaction level λ :

$$\text{Maximise } \lambda \text{ subject to } \mu(Z_1) \geq \lambda, \mu(Z_2) \geq \lambda, \text{ and constraints (i)–(iv)} \quad (11)$$

Table 4. Pay-off table (Kucukbay and Araz, 2016, Table 1)

Optimised goal	Expected return (%)	Risk (%)
Maximise return	4.059	16.326
Minimise risk	2.148	10.347

The portfolios were built on April 2009 – March 2014 data and held for the test year April 2014 – March 2015. They were judged with the Sharpe index $S = (r_p - r_f) / \sigma_p$, using a risk-free rate of $r_f = 0.722\%$, and compared with the BIST 100 index. The paper does not say which software solved the LPs.

3.3 Assumptions common to both studies

Both studies assume that past returns are a fair guide to the near future, that short selling and transaction costs can be ignored, and that money can be split into any amounts. Study 1 adds that the quadratic model needs normally distributed returns and a risk-averse investor, which the LP models do not require. Study 2 adds that its target ranges come from the decision maker's judgement and the pay-off table, and that only stocks traded throughout the period are considered.

4. RESULTS

4.1 Study 1

Table 5 shows the portfolios chosen from the twelve months of data. The quadratic model spread money over four shares and had the lowest return and risk. Maximin put everything into two shares and had the highest return and risk. MAD was close to maximin in return and risk here, but closer to the quadratic model in how widely it spread the money across shares.

Table 5. Optimal portfolios on 12-month data (Papahristodoulou, 2002, Table 3)

Model	Portfolio (share of budget)	Return $E(r)$	Risk σ
QP (Markowitz)	B 16.4%, C 33.1%, D 18.8%, E 31.6%	3.412%	1.847%
Maximin LP	C 45.96%, E 54.04% ($Z = 98.5$ SEK)	3.516%	2.286%
MAD LP	B 8.6%, C 36.3%, D 1.6%, E 53.5%	3.510%	2.260%

The author describes MAD as close to the quadratic model. On the figures above, however, MAD's risk (2.260%) is closer to maximin's (2.286%) than to the quadratic model's (1.847%). The claim therefore holds better for the shape of the portfolio than for its risk level, and we flag this as a point where the paper's summary is more generous than its table.

With $w = 1$ the utilities are 1.565 (QP), 0.074 (maximin) and 1.250 (MAD). When the investor strongly dislikes risk, QP gives the best utility, and when the investor accepts risk (w close to 0), maximin and MAD do better. Table 6 gives the out-of-sample results.

Table 6. Out-of-sample performance, return / risk in % (Papahristodoulou, 2002)

Model	Next 3 months	Next 6 months
QP (Markowitz)	-0.302 / 0.901	-0.656 / 0.745
Maximin LP	0.079 / 1.720	-0.685 / 1.551
MAD LP	0.043 / 1.490	-0.651 / 1.345

In the first three months the QP portfolio did worst. Over six months all three lost money, with MAD and QP losing least.

4.2 Study 2

Table 7. Training-period results (Kucukbay and Araz, 2016, Table 5)

	LPP	FGP
Number of stocks selected	12	13
Expected return	3.247%	3.4292%
Risk	11.706	12.3025

In training, FGP gives a slightly higher return but also higher risk. Table 8 shows how the portfolios fared in the test year.

Table 8. Out-of-sample performance, April 2014 – March 2015 (Kucukbay and Araz, 2016, Table 6)

	BIST 100 index	LPP portfolio	FGP portfolio
Return	2.40%	4.33%	3.90%
Risk	6.48%	8.41%	8.96%
Sharpe index	0.26	0.43	0.36

Check: $LPP \text{ Sharpe index} = (4.33 - 0.722) \div 8.41 \approx 0.43$.

Both portfolios beat the market index, and LPP had the best Sharpe index. LPP earned more than FGP in the test year, even though it earned less in training.

5. DISCUSSION

5.1 Comparing the two studies

Table 9. Comparison of the two studies

Feature	Study 1	Study 2
Objectives	One at a time (risk or worst return)	Two together (return and risk)
Risk measure	Variance (QP), worst month (maximin), mean absolute deviation (MAD)	Sum of each stock's standard deviation
Practical rules	Budget, minimum return, 75% cap per share	Budget, 25% per sector, $\geq 50\%$ in BIST 50, 10% cap per stock
Data size	5 stocks, 12 months	89 stocks, 60 months
Software named	QSB+ (for QP only)	Not stated
Validation	Next 3 and 6 months	Next 12 months with Sharpe index
Difficulty	Easy to compute by hand	Needs a weight algorithm and software

Study 1 shows how a linear model is built and how a non-linear risk measure is linearised. Study 2 shows how such models cope with several goals and realistic rules on a real market. They address the same problem at different levels of complexity.

5.2 Strengths

Study 1 is small, transparent and uses real data, so every number can be checked. It compares three models on the same data and tests them on new data. Its LP models need no covariance matrix, and they can be extended to integer programs to handle minimum share lots or fixed costs. Study 2 uses realistic data and rules, handles two conflicting goals at once, lets the investor state flexible satisfaction levels, and validates the results on a separate test year against the market index.

5.3 Limitations

Study 1. Five shares and twelve months are a very small sample, and the author says more examples with larger data are needed. The 3- and 6-month tests are short, and since all three portfolios lost money over six months, no model is clearly proved better. Transaction costs are ignored and the LP software is not named.

Study 2. Risk is measured as $\sum \sigma_i X_i$, a weighted sum of individual standard deviations. This ignores how stocks move together, so it overstates the risk of a diversified portfolio. Only one test year is used, so the claim that LPP beats FGP rests on limited evidence, and LPP did earn a lower return than FGP in training. The target ranges are partly subjective and the weight algorithm needs a computer program. Including only stocks that traded throughout the period can make past results look better than they were, and transaction costs and software are not reported.

We also noticed some points in the paper that deserve a closer look. Read plainly, the BIST 50 rule says at least half the money should be in BIST 50 firms, but the stock tables suggest that the LPP portfolio holds about 18% and the FGP portfolio about 30% in such firms. The equation for that rule is written ambiguously, so this may be a presentation problem rather than a modelling error, but it could not be confirmed from the paper alone. The range boundaries in Table 3 also overlap slightly (14.688 and 14.680), and the reference list contains several duplicated entries.

5.4 Possible improvements

For Study 1: use more shares and a longer period, add transaction costs and minimum lots (integer LP), and test other minimum-return levels. For Study 2: include covariance or a downside-risk measure such as mean absolute deviation, test over several years and market conditions, and add criteria such as liquidity and dividends.

6. CONCLUSION

Portfolio allocation is, at heart, a question of how to share out a limited pot of money. Markowitz gave it a rigorous answer, but one that gets heavy when the list of stocks grows, which is why people looked for linear alternatives. We reviewed two such efforts. On five Swedish stocks, Papahristodoulou showed that a mean-absolute-deviation model can pick a portfolio of similar shape to the quadratic one without needing a covariance matrix, though its risk ended up nearer to the maximin model, which chases a higher return at higher risk. On 89 Turkish stocks, Kucukbay and Araz showed that linear physical programming can handle two goals and a set of practical rules at once, and in their test year it earned the best Sharpe index (0.43) against fuzzy goal programming (0.36) and the market (0.26).

We should be honest about how far this goes. The first study has five stocks and a six-month test in which every portfolio lost money. The second has one test year, and measures risk in a way that ignores how stocks move together. So these papers show that linear models are workable and useful, not that they beat the quadratic model or each other. What stood out most to us is how much of the real work lies in modelling choices, such as what to count as risk and what rules to impose, rather than in the solving itself.

A natural next step would be to run the same models on larger and more recent data, over several test periods, with transaction costs, minimum lot sizes and a risk measure that accounts for how stocks move together. Until that is done, the two studies are best read as a convincing argument that linear programming deserves a place in portfolio selection, and not as the last word on which model to use.

REFERENCES

- Konno, H. and Yamazaki, H. (1991). Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market. *Management Science*, 37(5), 519–531.
- Kucukbay, F. and Araz, C. (2016). Portfolio selection problem: a comparison of fuzzy goal programming and linear physical programming. *An International Journal of Optimization and Control: Theories & Applications (IJOCTA)*, 6(2), 121–128. <https://doi.org/10.11121/ijocta.01.2016.00284>
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Papahristodoulou, C. (2002). *Optimal portfolios using linear programming models* (working paper). Mälardalen University, Västerås, Sweden. Available at Optimization Online: <https://optimization-online.org/wp-content/uploads/2002/10/549.pdf>
- Papahristodoulou, C. and Dotzauer, E. (2004). Optimal portfolios using linear programming models. *Journal of the Operational Research Society*, 55(11), 1169–1177. <https://doi.org/10.1057/palgrave.jors.2601765>