

# Heuristic approaches to realistic portfolio optimisation

F. Busetti

*Andisa Securities, South Africa*

## Abstract

Previous research in portfolio optimisation has incorporated some real-world aspects but, to the best of our knowledge, none has incorporated all of them in a model applicable to actual real-world portfolios.

We therefore develop a realistic model, investigate the efficiency of its solution by two heuristic methods, genetic algorithms and tabu search, and then examine the insights provided by the optimisation of real portfolios.

Our model is based on the classical mean-variance approach, enhanced with floor and ceiling constraints, cardinality constraints and nonlinear transaction costs that include a substantial illiquidity premium, and is applied to a large 100-stock portfolio.

We find that for large portfolios the performance of genetic algorithms is three orders of magnitude better than that of tabu search.

The results confirm that both floor and ceiling constraints have a substantial negative impact on real portfolio performance. Optimal portfolios with nonlinear costs and cardinality constraints often contain a large number of stocks with very low weightings. Their function is to diversify risk, and floor constraints hamper this, damaging portfolio performance. In addition, nonlinear transaction costs that are comparable in magnitude to forecast returns tend to diversify portfolios materially.

*Keywords: portfolio optimisation, efficient frontier, heuristic, genetic algorithm, tabu search.*

## 1 Introduction

While the basis for portfolio optimisation was established by Markowitz [1] in a seminal paper over 50 years ago, it is often difficult to incorporate real-world constraints into the classical theory. In practical portfolio construction there are



floor and ceiling constraints. In addition, the impact of transaction costs on performance can be large. These costs are variable and there is also an additional “liquidity premium” that drives up the cost for large transactions in illiquid stocks. This premium is strongly nonlinear and can be up to two orders of magnitude larger than the base costs. Combined with cardinality constraints this results in a mixed integer nonlinear programming problem that classical algorithms are unable to optimise efficiently.

## 2 Related work

Mulvey [2] approximated a transaction cost function by a piecewise linear convex function. Zenios and Kang [3] applied a linear programming model, using mean-absolute deviation (MAD) as the risk function, to a mortgage-backed securities portfolio. Konno and Wijayanayake [4] use a branch-and-bound algorithm to solve the MAD optimisation model for a concave cost function. Loraschi *et al.* [5] presented a distributed genetic algorithm for the unconstrained portfolio optimisation problem, while Crama and Schyns [6] developed a model incorporating many types of constraints but costs were ignored. The algorithms of Bienstock [7] exploit the fact that the objective function is quadratic and that the covariance matrix is of low rank, while Borchers and Mitchell [8] use an interior point nonlinear method. Mansini and Speranza [9] used a MAD approach and considered floor constraints. Chang *et al.* [10] constructed a cardinality-constrained Markowitz model incorporating floor and ceiling constraints that was solved using genetic algorithms, simulated annealing and tabu search, but costs were not addressed. Tabu search was applied by Glover *et al.* [11] to a portfolio optimisation problem. Lobo *et al.* [12] described an approximate method incorporating ceiling constraints, risk constraints and costs. Only linear and fixed transaction costs were used, and cardinality constraints were not incorporated.

## 3 The real-world model

The unconstrained Markowitz model is as follows. If:

$N$  = the number of assets in the investable universe

$R_i$  = the expected return of asset  $i$  ( $i = 1 \dots N$ ) above the risk-free rate  $r_f$

$\sigma_{ij}$  = the covariance between assets  $i$  and  $j$  ( $i = 1 \dots N, j = 1 \dots N$ )

$x_i$  = the fractional weight in the portfolio of asset  $i$  ( $i = 1 \dots N$ )

$w$  = the risk-aversion parameter ( $0 \leq w \leq 1$ )

then the problem becomes:

$$\text{maximise} \quad (1-w) \sum_{i=1}^N R_i x_i - w \sum_{i=1}^N \sum_{j=1}^N \sigma_{ij} x_i x_j \quad (1)$$

$$\text{subject to} \quad \sum_{i=1}^N x_i = 1$$

$$0 \leq x_i \leq 1 \quad i = 1, \dots, N$$



Solving the quadratic programming equation (1) for various values of  $w$  results in combinations of portfolio return and variance that map the efficient frontier. This curve represents the set of Pareto-optimal or non-dominated portfolios. Our real-world extensions to the above model are as follows.

**3.1 Floor and ceiling constraints**

If  $l_i$  = the minimum weighting that can be held of asset  $i$  ( $i = 1 \dots N$ )  
 $u_i$  = the maximum weighting that can be held of asset  $i$  ( $i = 1 \dots N$ )  
 then the constraint is simply formulated as

$$l_i \leq x_i \leq u_i \tag{2}$$

where  $0 \leq l_i \leq u_i \leq 1$  ( $i = 1 \dots N$ )

**3.2 The cost function**

The conceptual shape of the transaction cost function is shown in figure 1, where units can refer to either number of stocks or transaction size in monetary units.

Our model is restricted to the high end of the cost curve, as this is the region relevant to institutional investors. If:

- $m$  = marketable securities tax (MST) rate
- $f$  = a fixed charge component
- $v$  = value-added tax (VAT) rate
- $b$  = brokerage rate
- $s$  = transaction value
- $p$  = the illiquidity premium (a function)
- $c'$  = total unit transaction cost
- $C$  = total transaction cost

then the total transaction cost is given by

$$C = (1 + v)f + [(1 + v)(b + p) + m]s \tag{3}$$

and the total unit transaction cost is

$$c' = C/s = (1 + v)f/s + [(1 + v)(b + p) + m] \tag{4}$$

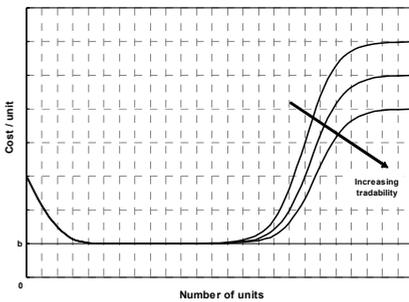


Figure 1: Illustrative transaction cost functions.

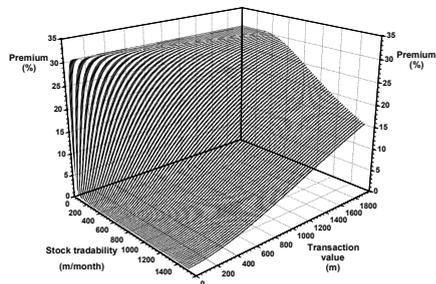


Figure 2: Illiquidity premium surface.



The illiquidity premium is mainly a function of transaction size relative to the stocks' tradability. The illiquidity premium is therefore given by a function of  $s/t$ , where  $t$  = stock tradability (average value traded on the market per time period). Our estimates of the illiquidity premium  $p$  are fitted empirically with a two-term exponential function of the form

$$p(s/t) = a[1 - ke^{-\{(k-1)/k\} d (s/t)} + (k-1)e^{-d (s/t)}] \quad (5)$$

using the three parameters  $a$ ,  $k$  and  $d$ .

The resulting set of curves of the illiquidity premium for various values of  $t$  is shown in figure 2. As stock tradability increases the cost curve declines and also becomes more linear.

### 3.3 Cardinality constraint

Let  $z_i = 1$  if any amount of asset  $i$  ( $i = 1, \dots, N$ ) is held  
 $z_i = 0$  otherwise

$K$  = the maximum number of assets allowed in the portfolio

Then the cardinality constraint becomes

$$\sum_{i=1}^N z_i = K \quad (6)$$

where  $K \leq N$  and  $z_i \in [0, 1]$  is the integrality constraint.

## 4 Solving the real-world model

### 4.1 The heuristic approach

Heuristic solution methods are approximate algorithms that are applied to intractable problems and ensure ever-improving solutions that will result in "acceptable" answers. Examples of heuristic algorithms include simulated annealing, tabu search and genetic algorithms. Since it is difficult to find commercial software for simulated annealing, we test tabu and genetic algorithms on the problem.

### 4.2 Cardinality-unconstrained case

The cardinality-unconstrained case is first examined to establish the impact of the real-world features of this model. A small portfolio of 10 stocks is sufficient for this purpose and is easily optimised with a simple solver. The 10 largest stocks by market capitalisation are used in this analysis, with a floor constraint of 2% and ceiling constraint of 15%.

#### 4.2.1 Effect of floor and ceiling constraints

The effect of floor and ceiling constraints on the efficient frontier is shown in figure 3. Risk is shown as the standard deviation returns. Floor constraints force an exposure to every stock, including those with poor returns, thus reducing the



portfolio's return. For low levels of risk aversion, the portfolio will normally tend to consist of only one or two stocks, i.e. those with the highest returns.

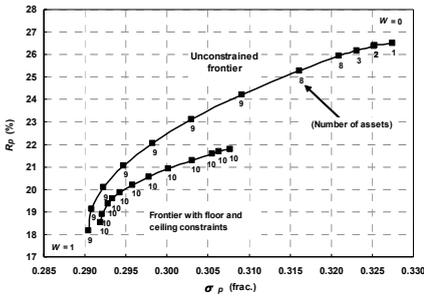


Figure 3: Effect of floor and ceiling constraints (no costs).

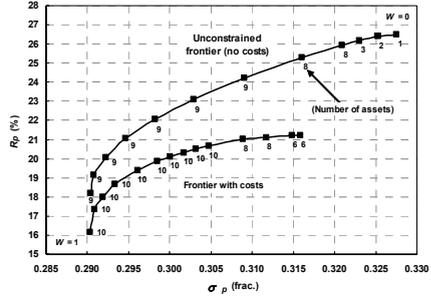


Figure 4: Effect of costs (no floor and ceiling constraints).

The ceiling constraint, however, will make this high level of optimal exposure to these stocks impossible, and again force an exposure to lower-return stocks, which will reduce the portfolio's return.

The lowest-risk constrained portfolio has a higher volatility than that of the unconstrained portfolio, since the constraints also interfere with the optimal weights for risk reduction.

#### 4.2.2 Effect of nonlinear costs

The impact of nonlinear costs on the portfolio, without any floor or ceiling constraints, is shown in figure 4.

Without costs, the highest-returning portfolio again consists only of one stock. However, the large size of the transaction results in a high transaction cost since costs increase exponentially with transaction size because of the illiquidity premium. If this cost is of a magnitude comparable to the forecast returns, the portfolio tends to diversify into more stocks in order to reduce total transaction costs and their adverse impact on returns.

Therefore, realistic nonlinear transaction costs tend to diversify portfolios.

#### 4.2.3 Combined effect of floor and ceiling constraints and nonlinear costs

The impact on the portfolio of floor and ceiling constraints together with nonlinear costs is shown in figure 5.

The negative impact on the constrained and cost-laden portfolio is cumulative. The characteristics of the three realistically-constrained frontiers are summarised in table 1.

Risk for the highest-return portfolio is reduced by the introduction of constraints and nonlinear costs, since they diversify the portfolio. However, for the lowest-risk portfolio, floor and ceiling constraints will increase risk since they force exposure to higher-risk stocks. An interesting result is that this may in some cases also be accompanied by a corresponding increase in return. The



impact of nonlinear costs on risk is thus ambiguous, depending on the degree of diversification required for cost reduction. The number of stocks in the portfolio invariably increases as a result of constraints (trivially), nonlinear costs and their combined effect.

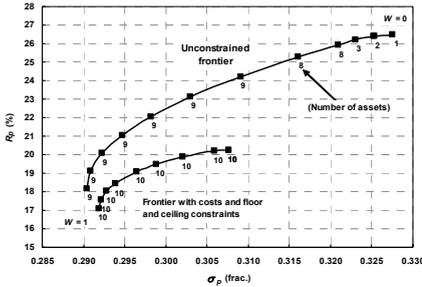


Figure 5: Combined effect of constraints and costs.

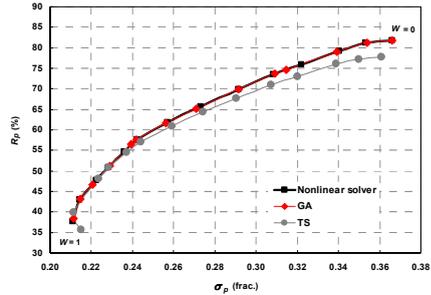


Figure 6: Heuristic-derived 100-stock cardinality-unconstrained frontier.

Table 1: Summary of constraint and cost effects.

Model	Portfolio	Highest-return portfolio	Lowest-risk portfolio
No constraints	Return (%)	26.50	18.18
	Risk (frac)	0.328	0.290
	Stocks no.	1	8
With floor and ceiling constraints	Return (%)	21.78	18.53
	Risk (frac)	0.308	0.292
	Stocks no.	10	10
With nonlinear costs	Return (%)	21.24	16.16
	Risk (frac)	0.316	0.290
	Stocks no.	6	9
With floor and ceiling constraints and nonlinear costs	Return (%)	20.24	17.07
	Risk (frac)	0.308	0.292
	Stocks no.	10	10
Difference between all constraints and no constraints	Return (%)	-6.26	-1.11
	Risk (frac)	-0.020	0.002
	Stocks no.	9	2

### 4.3 Cardinality-constrained case

#### 4.3.1 Testing the heuristic methods

The incorporation of cardinality constraints is essential in any realistic portfolio optimisation.

For the real-world cardinality-constrained, large 100-stock portfolio with both floor and ceiling constraints and nonlinear costs which we now consider, there is



no classical method of calculating the exact efficient frontier because of the size and nonlinearity of the problem, and hence no way of benchmarking the heuristic methods against the exact solution. Therefore, to test initially the effectiveness of the heuristic methods and establish their suitability, they are first used to find the efficient frontier without the cardinality constraints.

#### 4.3.2 Evaluating the heuristics on the cardinality-unconstrained case

The efficient frontiers generated by the two heuristic methods for the cardinality-unconstrained but large 100-stock portfolio with floor/ceiling constraints and nonlinear costs are shown in figure 6.

The comparative results of genetic algorithms (GA) and tabu search (TS) are summarised in table 2. The closeness of the calculated frontier to the exact frontier is measured by the arithmetic average of the absolute percentage differences in both return and risk. Solution times are for a 2.8 GHz computer.

Table 2: Performance of heuristics.

Absolute difference in:		Return (%)	Risk (%)	Objective function (%)	Solution time (min)	Best trial (no.)	Total trials (no.)
TS	Median	3.30	0.54	8.56	30	371	723
	Standard deviation	2.06	0.46	20.62	3	271	175
	Mean	3.06	0.73	16.48	29	379	736
	Combined mean	1.89		-	-	-	-
GA	Median	0.36	0.24	0.07	4	14 819	14 819
	Standard deviation	0.95	0.62	0.09	4	10 656	10 656
	Mean	0.85	0.51	0.08	5	18 419	18 419
	Combined mean	0.68		-	-	-	-

While tabu search works well on small (around 20-stock) problems, its rate of convergence to the solution slows significantly when the problem size increases towards 100 stocks. A stopping rule of 30 minutes is therefore implemented.

A simple efficiency measure is provided by combining absolute accuracy and time by using their product as the criterion. Figure 7 shows both methods' performance across the frontier for various values of  $w$ .

Based on this measure, the performance of the genetic algorithm is better than that of tabu search by approximately three orders of magnitude and is able to find solutions arbitrarily close to the correct value, with calculation times of around 5 minutes.

Interestingly, both methods find the centre portion of the efficient frontier the most difficult to generate. A possible reason is that at the upper end (highest returns, low risk aversion,  $w \Rightarrow 0$ ), the selection of the highest-return stocks is relatively straightforward, and at the lower end (lowest risk, high risk-aversion,  $w \Rightarrow 1$ ) the optimisation strategy is also simple: select the lowest-risk stocks. However, in the central part of the frontier there is a much larger number of

combinations of stocks that will result in middle-of-the-road return and risk levels.

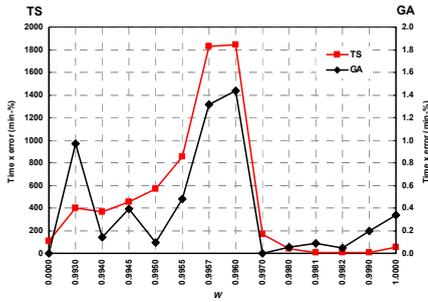


Figure 7: Efficiency of heuristic methods.

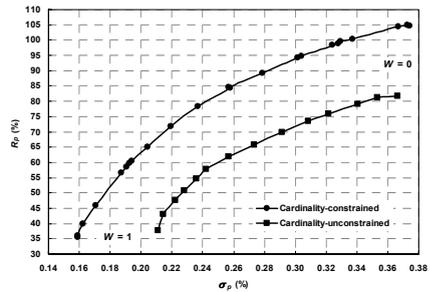


Figure 8: Cardinality-constrained efficient frontier.

Since the efficient frontier found by the heuristic methods consists of suboptimal points, the optimal portfolios generated for the cardinality-constrained case will always be conservative in terms of displaying returns per unit of risk that are lower than their true values.

### 4.3.3 Application to the cardinality-constrained case

In the cardinality-constrained case the floor constraint is subsumed into the cardinality count, i.e. if  $x_i > l_i$  then  $z_i = 1$ , otherwise it is zero. The ceiling of 15% is retained. A cardinality constraint of 40 stocks within the 100-stock universe is selected.

The introduction of cardinality constraints may result in a discontinuous efficient frontier, as shown by Chang *et al* [10]. The efficient frontier for 40 stocks is shown in figure 8. There are no signs of any discontinuities in this particular cardinality-constrained efficient frontier.

The cardinality-constrained portfolio completely dominates the cardinality-unconstrained portfolio. On average, for the same level of return the cardinality-constrained frontier exhibits risk that is lower by between 5% and 12%. Conversely, for equal risk levels, the cardinality-constrained portfolio produces returns that are 24% to 30% higher across the efficient frontier.

The next step in optimising real-world portfolios is to determine the risk-aversion factor  $w$ . This can be estimated from the original Markowitz theory, i.e. by drawing the capital market line from the point representing risk-free T-bills to the efficient frontier and noting the value of  $w$  at the point of tangency.

Using this value of  $w$ , optimal 40-stock portfolios are then generated from the 100-stock universe. Their characteristics are compared with those of the universe in table 3, which also shows the effect of changing the floor constraint from 0.5% to 2%. Since the higher floor constraint shifts the frontier, a different value of  $w$  arises.



Table 3: Optimal portfolio characteristics.

	Weight range $x_i$ (frac)	Total costs $C(x)$ (%)	Excess return less costs $R_i$ (%)	Risk (beta) (x)
100-stock universe				
Maximum return	0.150	1.72	146.5	1.81
Equally-weighted	0.010	0.29	40.2	1.11
Minimum risk	0.000	0.00	-10.4	0.26
Optimal portfolio	-	0.60	39.9	1.11
40-stock optimal portfolios				
Parameters:		Floor =	0.005	$w = 0.996$
Maximum return	0.150	1.72	146.5	1.61
Equally-weighted	0.025	0.73	59.5	1.00
Minimum risk	0.005	0.59	17.5	0.26
Optimal portfolio	-	1.28	83.5	0.89
Parameters:		Floor =	0.020	$w = 0.998$
Maximum return	0.150	1.72	146.8	1.61
Equally-weighted	0.025	0.65	56.5	0.99
Minimum risk	0.020	0.62	-4.6	0.26
Optimal portfolio	-	0.79	60.1	0.90

The substantially better performance of the 40-stock cardinality-constrained portfolios in terms of providing substantially higher returns at lower risk is readily apparent. The optimal portfolios generated usually consists of relatively few stocks with high weights that are at, or close to, the ceiling constraint, a larger but still relatively small number of medium weightings and a long tail consisting of many stocks at, or close to, the floor constraint. In terms of the number of stocks, this tail can be around 70%-80% of the portfolio. What happens is that the highest weightings are usually allocated to stocks with high forecast returns. However, these stocks normally also have above-average risk, which raises the portfolio's risk level. This risk is then diversified away by the large number of stocks with very low weightings.

## 5 Conclusions

We have shown that realistically large portfolios, incorporating floor and ceiling constraints, nonlinear transaction costs including a substantial illiquidity premium, together with cardinality constraints, can be optimised in reasonably short times using genetic algorithms.

It should be noted that the addition of embellishments to the model, such as setting individual floor and ceiling constraints for each stock; applying other constraints such as market capitalisation or tradability, as well as class, sector or style constraints; providing individual cost curves for individual stocks; setting the cardinality constraint as a range and constructing more complex objective functions and cost curves (which could incorporate step functions or even be discontinuous) will be equally easily handled by this approach.



## References

- [1] Markowitz, H., Portfolio selection. *Journal of Finance*, **7**, pp.77-91, 1952.
- [2] Mulvey, J. M., Incorporating transaction in models for asset allocation. *Financial Optimization*, ed. W. Ziemba et al, Cambridge University Press, pp. 243-259, 1993.
- [3] Zenios, S. A. and Kang, P., Mean-absolute deviation portfolio optimization for mortgage-backed securities. *Annals of Operations Research*, **45**, pp. 433-450, 1993.
- [4] Konno, H. and Wijayanayake, A., Mean-absolute deviation portfolio optimization model under transaction costs. *Journal of the Operations Research Society of Japan*, **42**, pp. 422-435, 1999.
- [5] Loraschi, A., Tettamanzi, A., Tomassini, M. and Verda, P., Distributed genetic algorithms with an application to portfolio selection problems. *Artificial Neural Nets And Genetic Algorithms*, ed. D. W. Pearson, N.C. Steele and R. F. Albrecht, Springer-Verlag, pp. 384-387, 1995.
- [6] Crama, Y. and Schyns, M., Simulated annealing for complex portfolio selection problems. *European Journal of Operational Research*, **150**, pp. 546-571, 2003.
- [7] Bienstock, D., Computational study of a family of mixed-integer quadratic programming problems. *Mathematical Programming*, **74**, pp. 121-140, 1996.
- [8] Borchers, B. and Mitchell, J. E., A computational comparison of branch and bound and outer approximation algorithms for 0-1 mixed integer nonlinear programs. *Computers & Operations Research*, **24**, pp. 699-701, 1997.
- [9] Mansini, R. and Speranza, M. G. Heuristic algorithms for the portfolio selection problem with minimum transaction lots. *European Journal of Operational Research*, **114**, pp. 219-233, 1999.
- [10] Chang, T.-J., Meade, N., Beasley, J. E. and Sharaiha, Y. M., Heuristics for cardinality constrained portfolio optimization. *Computers & Operations Research*, **27(13)**, pp. 1271-1302, 2000.
- [11] Glover, F., Mulvey, J. M. and Hoyland, K., Solving dynamic stochastic control problems in finance using tabu search with variable scaling. *Metaheuristics: theory & applications*, ed. I. H. Osman and J.P. Kelly, Kluwer Academic Publishers, pp. 429-448, 1996.
- [12] Lobo, M. S., Fazel, M. and Boyd, S., Portfolio optimisation with linear and fixed transaction costs and bounds on risk. (Revise-and-resubmit decision from *Annals of Operations Research*, January 2005.)

