On the Effectiveness of a NSGA-II Local Search Approach Customized for Portfolio Optimization

Kalyanmoy Deb[a], Ralph Steuer[b], Rajat Tewari[c] and Rahul Tewari[d]

[a] Indian Institute of Technology Kanpur, Kanpur, PIN 208016, India, deb@iitk.ac.in,
Also Department of Information and Service Economy, Aalto University, School of Economics, 00100 Helsinki, Finland

[b] Department Banking and Finance, Terry College of Business, University of Georgia, Athens, Georgia 30602-6253, USA, rsteuer@uga.edu

[c] Benaras Hindu University, Benaras, India, rajat.tewari.min08@itbhu.ac.in

[d] Deutsche Bank Group, Mumbai, India, rtewari.iitk@gmail.com

KanGAL Report Number 2011007

Abstract

Bi-objective portfolio optimization for minimizing risk and maximizing expected return has received considerable attention using evolutionary algorithms. Although the problem is a quadratic programming (QP) problem, the practicalities of investment often make the decision variables discontinuous and introduce other complexities. In such circumstances, usual QP solution methodologies can not always find acceptable solutions. In this paper, we modify a bi-objective evolutionary algorithm (NSGA-II) to develop a customized hybrid NSGA-II procedure for handling situations that are non-conventional for classical QP approaches. By considering large-scale problems, we demonstrate how evolutionary algorithms enable the proposed procedure to find fronts, or portions of fronts, that can be difficult for other methods to obtain. In addition, post-optimality analyses are performed to reveal salient properties of optimal solutions that can remain as vital knowledge to a practitioner.

1 Introduction

Portfolio optimization inherently involves conflicting criteria. Among possible objectives, minimizing risk and maximizing expected return (also known as the mean-variance model of Markowitz [17]) are the two objectives that have received the most attention. The decision variables in these problems are the proportions of initial capital to be allocated to the different available securities. Such bi-objective problems give rise to fronts of trade-off solutions which must be found to investigate the risk-return relationships in a problem.

In addition to the two objectives, these problems possess constraints. While the calculation of expected return is a linear function of the decision variables, the calculation of portfolio risk is a quadratic function of the decision variables. Thus, the overall problem, in its simplest form, is a bi-objective quadratic programming (QP) problem, and when all constraints are linear and all variables are continuous, the fronts of such problems can be solved for exactly using the QP solvers [21, 18]. However, in practice, there can be conditions which make QP solvers difficult to apply. For instance, a portfolio with very small investments in one or more securities is likely not to be of interest on managerial grounds alone. This creates a need for decision variables that either take on a value of zero or a non-zero value corresponding to at least a minimum investment amount. Moreover, users may only be interested in portfolios that involve limited numbers of securities. In addition, there can be governmental restrictions contributing to the complexity of the process, thus only adding to the difficulties of handling portfolio problems by classical means.

Evolutionary multi-objective optimization (EMO) has been found to be useful in solving many different optimization problems. EMO has also been experimented with in portfolio optimization problems [21, 2, 25, 16, 15]. For example, Chang et al. [3] used genetic algorithms (GAs), tabu search, and simulated annealing on portfolio problem with given cardinalities on the number of assets, but the problems on which the approaches were tested were small. Other approaches including simulated annealing [4], differential evolution [14], and local search based memetic algorithms [23, 22] have also been attempted. Here, we suggest and develop a customized hybrid NSGA-II procedure which is uniquely designed to handle the non-smooth conditions that can often appear in portfolio analysis. The most important aspect of the proposed algorithm is that the initialization procedure and the recombination and mutation operators are all customized so that the proposed procedure starts with a feasible population and *always* creates only feasible solutions. Conflicting objectives are handled using the elitist non-dominated sorting GA or NSGA-II [8]. To make the obtained solutions close to optimal solutions, the NSGA-II solutions are clustered into small groups and then a local search procedure is commenced from each solution until no further improvements are possible.

Besides developing a customized NSGA-II for the portfolio optimization problems, the obtained trade-off solutions are analyzed to find important investment patterns, if any, common to the solutions. Certain important and generic knowledge about investments relating to minimizing risk and maximizing return are revealed.

The remainder of the paper is organized as follows. Section 2 discusses portfolio optimization in greater detail. Thereafter, the customized hybrid NSGA-II procedure for solving portfolio optimization problems is described in Section 3. The local search component of the overall approach is described next. Section 4 presents results obtained by the proposed procedure. Results are analyzed to reveal important investment patterns. Concluding remarks constitute Section 5.

2 Practical Portfolio Optimization

In a portfolio problem with an asset universe of n securities, let x_i ($i = 1, 2, \ldots, n$) designate the proportion of initial capital to be allocated to security i. And typically there are two conflicting goals – to minimize portfolio risk while attempting to maximize portfolio return [17]. In its most basic form, this results in the following problem:

$$\begin{array}{rl} \text{Minimize} & f_1(\mathbf{x}) = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i \sigma_{ij} x_j, \\ \text{Maximize} & f_2(\mathbf{x}) = \sum_{i=1}^{n} r_i x_i, \\ \text{subject to} & \sum_{i=1}^{n} x_i = 1, \\ & x_i \geq 0. \end{array} \tag{1}$$

The first objective is portfolio risk that is usually computed from an $n \times n$ covariance matrix $[\sigma_{ij}]$. The second objective is expected portfolio return as computed from a weighted sum of the individual security expected returns. The first constraint ensures the investment of all funds. The second constraint ensures the non-negativity of each investment.

With the objectives conflicting, the solution to the above is the set of all of the problem's Pareto-optimal solutions as this is precisely the set of all of the problem's contenders for optimality. One of the ways to solve for the front of (1) is to convert the problem into a number of single-objective problems, the popular way being to subordinate the expected return objective to a

$\geq$-type constraint as in the ϵ-constraint [20] formulation:

$$\begin{array}{rl} \text{Minimize} & f_1(\mathbf{x}) = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i \sigma_{ij} x_j, \\ \text{subject to} & \sum_{i=1}^{n} r_i x_i \geq R, \\ & \sum_{i=1}^{n} x_i = 1, \\ & x_i \geq 0. \end{array} \tag{2}$$

The above QP problem is solved repetitively for many different values of R to generate a non-dominated front.

Other than in small problems, it can be expected that almost any solution of (2) will contain many of its securities at the zero level. That is, it can be expected that for many i, $x_i^* = 0$. It can also be expected for at least a few securities that x_i^* will be a very small number. However, to have a practical portfolio, very small investments in any security may not be desired and are to be avoided. Thus, there is the practicality that for any portfolio to be of interest, there is to be a lower limit on any non-zero investment. That is, either $x_i^* = 0$ (meaning no investment in the i-th security) or $x_i^* \geq \alpha$ (meaning that there is a minimum non-zero investment amount for the i-th security). There may also be an upper bound β on the proportion of any security in any portfolio. Unfortunately, the solution of (2) for any given R does not guarantee the possession of any of these characteristics.

In addition to the above, there is a second practicality and it is about the number of non-zero securities contained in the portfolios along the Pareto front. Over this, a user may wish to exert control. To generate practical portfolios, a user may be interested in specifying a given number of non-zero investments or a range in the number of non-zero investments a portfolio is to contain. This is a cardinality constraint and it has also been the subject of some research attention [18, 22]. Taking both practicalities into account, we have the following bi-objective optimization problem:

$$\begin{array}{ll} \text{Minimize} & f_1(\mathbf{x}) = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i \sigma_{ij} x_j, \\ \text{Maximize} & f_2(\mathbf{x}) = \sum_{i=1}^{n} r_i x_i, \\ \text{subject to} & \\ \text{(1st constraint)} & \sum_{i=1}^{n} x_i = 1, \\ \text{(2nd constraint)} & x_i = 0 \text{ or } \alpha \leq x_i \leq \beta, \\ \text{(3rd constraint)} & d^{\min} \leq d(\mathbf{x}) \leq d^{\max}, \end{array} \tag{3}$$

where $\alpha > 0$ and $d(\mathbf{x})$ is given as follows:

$$d(\mathbf{x}) = \sum_{i=1}^{n} \left\{ \begin{array}{ll} 1, & \text{if } x_i > 0, \\ 0, & \text{if } x_i = 0. \end{array} \right. \tag{4}$$

2.1 Difficulties with Classical Methods

Standard QP solvers face difficulties in the presence of discontinuities and other complexities. For instance, the 2nd constraint requires an 'or' operation. While $x_i = 0$ or $x_i = \alpha$ is allowed, values between the two are not. This introduces discontinuities in the search space. The 3rd constraint involves a parameter d which is defined by a discontinuous function of the decision variables, given in (4). It is the presence of the 2nd and 3rd constraints that makes the application of standard methodologies difficult. In the following section, we discuss the GA based methodology of this paper for dealing with constraints like these.

3 Customized Hybrid NSGA-II Procedure

Generic procedures for handling constraints in an evolutionary algorithm exist [5, 12]. They do not care about the structure of the constraints, however they may not be efficient in handling a particular type of constraint. Here, we suggest a customized hybrid NSGA-II procedure for handling the specific constraints of the portfolio optimization problem.

There are three constraints in this problem. The first requires all variables to sum to one. Most evolutionary based portfolio optimization methodologies suggest the use of random keys or a dummy variable (x_i) and repair the variable vector as $x_i \leftarrow x_i / \sum_{i=1}^{n} x_i$ to ensure satisfaction of this constraint. In the presence of the first constraint alone, this strategy of repairing a solution is a good approach, but when other constraints are present, this may not be an adequate strategy.

The second constraint introduces an 'or' operation between two constraints representing two disjointed feasible regions. One approach suggested in the literature is to use an additional Boolean variable $\rho_i \in [0, 1]$ for each decision variable x_i, as follows [19]:

$$\alpha\rho_i \leq x_i \leq \beta\rho_i.$$

When $\rho_i = 0$, the variable $x_i = 0$ (meaning no investment in the i-th security). But when $\rho_i = 1$, the variable x_i takes a value in the range $[\alpha, \beta]$. This is an excellent fix-up of the 'or' constraint, but it introduces n new Boolean variables into the problem.

The third constraint involves a discontinuous function of decision variables. If the Boolean variables are used to take care of the set of second constraint, the cardinality constraint can be replaced by the following:

$$d^{\min} \leq \sum_{i=1}^{n} \rho_i \leq d^{\max}.$$

The presence of equality and inequality constraints would make it difficult for the generic penalty function approach to find a feasible solution and solve the problem.

Here, we suggest a different approach which does not add any new variable, but instead considers all constraints in its coding and evaluation procedures. First, a customized initialization procedure is used to ensure creation of feasible solutions. Thereafter, we suggest recombination and mutation operators which also ensure the creation of feasible solutions.

3.1 Customized Initialization Procedure

To create a feasible solution, we randomly select an integer value for d from the interval $[d^{\min}, d^{\max}]$. Then we create a portfolio which has exactly d non-zero securities. That is, we randomly select d of the n variables and assign zeros to the other $(n - d)$ variables. Since the non-zero x_i values must lie within $[\alpha, \beta]$, we select a random value from within the range of each of the non-zero variables. However, the random values of the non-zero variables may not sum to one, thereby violating the first constraint. To make the solution feasible, we repair the solution as follows. We map the non-zero variables to another set of non-zero variables in the range $[\alpha, \beta]$ in such a way that they sum to one. The mapping is illustrated in Figure 1.
If $\sum_{i=1}^{n} x_i = 1$, we have a feasible solution, otherwise we proceed as follows.

1. If $\sum_{i=1}^{n} x_i \neq 1$, the decision vector $\mathbf{x}$ is modified as follows:
   - For each i, if x_i is within $[\alpha, \alpha + \Delta_1]$, set $x_i = \alpha$.
   - For each i, if x_i is within $[\beta - \Delta_2, \beta]$, set $x_i = \beta$.
   - For each i, if x_i is within $[\alpha + \Delta_1, \beta - \Delta_2]$, x_i is kept the same.
2. Compute $X = \sum_{i=1}^{n} x_i$. If X is not one, then stretch the mapping line by moving its mid-point by an amount δ, as follows:
   - If $X < 1$, move the mid-point upwards by a distance δ in a direction perpendicular to the line (as shown in the Figure 1). This process is continued until $\sum_{i=1}^{n} x_i$ is arbitrarily close to one.

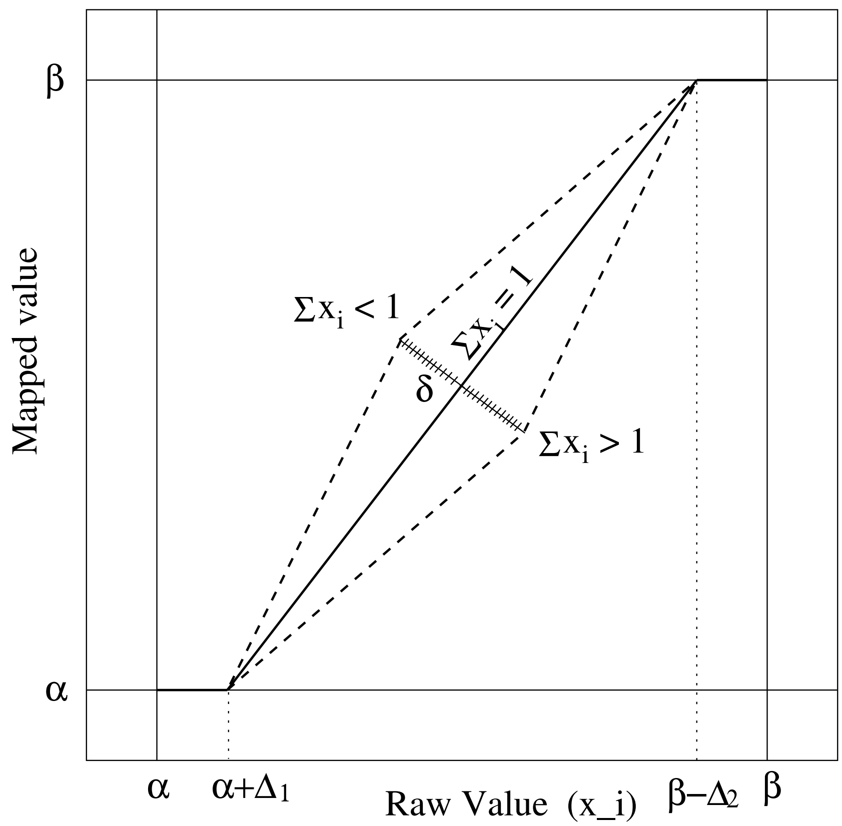

Figure 1: Mapping to satisfy the first constraint.

- If $X > 1$, move the mid-point downwards by a distance δ in a direction perpendicular to the line. This process is continued until $\sum_{i=1}^{n} x_i$ is arbitrarily close to one.

The parameters Δ_1 and Δ_2 allow us to map near-boundary values to their boundary values. However, if this is unwarranted, both of these parameters can be set to zero, particularly when all x_i values are either within $[\alpha, \alpha+\Delta_1]$ or within $[\beta-\Delta_2, \beta]$. The parameter δ may be chosen to be a small value and it helps to iterate the variable values so that the 1st constraint is satisfied within a small allowable error bound. A small value of δ may require a large number of iterations to satisfy the 1st constraint but the difference between $\sum_{i=1}^{n} x_i$ and one will be small, and vice versa. It remains as an interesting study to establish a relationship between δ and the corresponding error in satisfying the equality constraint, but in all of our simulations, we have chosen the following values for the parameters: $\Delta_1 = 0$, $\Delta_2 = 0.001$ and $\delta = 10^{-6}$.

After this repair mechanism, the 1st constraint is expected to be satisfied within a small tolerance. The 2nd and 3rd constraints are also satisfied in the process. Thus, every solution created by the above process is expected to be feasible.

3.2 Customized Recombination Operator

Once a feasible initial population is created, it is to be evaluated and then recombination and mutation operators are to be applied. Here, we suggest a recombination operator which always produces feasible offspring solutions by recombining two feasible parent solutions.

Two solutions are picked at random from the population as parents. Let the number of non-zero securities in the parents be d_1 and d_2, respectively. We illustrate the recombination procedure through the following solutions having $n = 10$:

$$\begin{array}{lcccccccccc} \text{Parent1:} & a_1 & 0 & a_3 & 0 & a_5 & 0 & 0 & a_8 & 0 & a_{10} \\ \text{Parent2:} & b_1 & 0 & b_3 & b_4 & b_5 & 0 & b_7 & 0 & 0 & b_{10} \end{array}$$

It is clear that $d_1 = 5$ and $d_2 = 6$ in the above. The following steps are taken to create a child solution.

1. A integer d_c is randomly selected from the range $[d_1, d_2]$. Let us assume that this value is $d_c = 5$, for our example.

2. The child solution inherits a zero investment for a particular security if both parents have zero investments in that security. For the above parents, this happens for three (n_0) securities: $i = 2$, 6 and 9. Thus, the partial child solution is as follows:

$$\begin{array}{lcccccccccc} \text{Child1:} & - & 0 & - & - & - & 0 & - & - & 0 & - \end{array}$$

3. The securities which have non-zero investments in both parents will inherit a non-zero investment value, but the exact amount will be determined by a real-parameter recombination operator (bounded form of SBX [7]) applied to the parent values. This operator ensures that values are not created outside the range $[\alpha, \beta]$. We discuss this procedure a little later. For the above parents, the child solution then takes on the form:

$$\begin{array}{lcccccccccc} \text{Child1:} & c_1 & 0 & c_3 & - & c_5 & 0 & - & - & 0 & c_{10} \end{array}$$

   Note that the number of common non-zero investments (n_c) between two parent solutions is such that $n_c \in [0, \min(d_1, d_2)]$. In our case of discussion here, $n_c = 4$.

4. The number of slots that remain to be filled with non-zero values is $w = d_c - n_c$. Since, $d_c \in [d_1, d_2]$, w is always greater than or equal to zero. From the remaining $(n - n_0 - n_c)$ securities, we choose w securities at random and take the non-zero value from the parent to which it belongs. For the above example parents, $w = 5 - 4 = 1$ and there are $(10 - 3 - 4)$ or 3 remaining places to choose the $w = 1$ security from. The remaining securities are assigned a zero value. Say, we choose the seventh (chosen at random) security to have a non-zero value. Since the non-zero value occurs in parent 2, the child inherits b_7. Thus, the child solution looks like the following:

$$\begin{array}{lcccccccccc} \text{Child1:} & c_1 & 0 & c_3 & 0 & c_5 & 0 & b_7 & 0 & 0 & c_{10} \end{array}$$

5. The above child solution is guaranteed to satisfy the 2nd and 3rd constraints, but it also needs to satisfy the 1st constraint. We use the procedure described in subsection 3.1 to repair the solution to a feasible one.

We also create a second child from the same pair of parents using the same procedure. Due to the creation of random integer d_c and other operations involving random assignments, the second child is expected to be different from the first child. Nevertheless, both child solutions are guaranteed to be feasible.

The bounded form of the SBX operator is described here for two non-zero investment proportions (say, a_1 and b_1, taken from parents 1 and 2, respectively). The operator requires a user-defined parameter η_c (usually set in the range [5,20] for multi-objective optimization [9]).

1. Randomly select a number u_i from $[0, 1]$.

2. Calculate γ_{qi} using the equation:

$$\gamma_{qi} = \begin{cases} (\kappa u_i)^{\frac{1}{\eta_c+1}}, & \text{if } u_i \leq \frac{1}{\kappa}, \\ \left(\frac{1}{2-\kappa u_i}\right)^{\frac{1}{\eta_c+1}}, & \text{otherwise.} \end{cases} \tag{5}$$

For $a_1 \leq b_1$, $\kappa = 2 - \zeta^{-(\eta_c+1)}$ and ζ is calculated as follows: $\zeta = 1 + 2\min[(a_1 - \alpha),\ (\beta - b_1)]/(b_1 - a_1)$. For $a_1 > b_1$, the role of a_1 and b_1 can be interchanged and the above equation can be used.

3. Then, compute the child from the following:

$$c_1 = 0.5[(1 + \gamma_{qi})a_1 + (1 - \gamma_{qi})b_1]. \tag{6}$$

The above procedure allows a zero probability of creating any child solution outside the prescribed range $[\alpha, \beta]$ for any two solutions a_1 and b_1 within the same range.

3.3 Customized Mutation Operator

Mutation operators perturb the variable values of a single population member. We employ two mutation operators here.

3.3.1 Mutation 1

In this operation, a non-zero security value is perturbed within its neighborhood by using the polynomial mutation operator [6]. For a particular value a_1, the following procedure is used along with a user-defined parameter η_m (usually in [20,50]):

1. Randomly select a number u_i from $[0, 1]$.

2. Calculate the parameter $\bar{\mu}$ as follows:

$$\bar{\mu} = \begin{cases} (2u)^{\frac{1}{\eta_m+1}} - 1, & \text{if } u \leq 0.5, \\ 1 - [2(1-u)]^{\frac{1}{\eta_m+1}}, & \text{otherwise.} \end{cases} \tag{7}$$

3. Calculate the mutated value, as follows:

$$a_1' = a_1 + \bar{\mu}\min[(a_1 - \alpha),\ (\beta - a_1)].$$

The above procedure ensures that a_1' lies within $[\alpha, \beta]$ and values close to a_1 are preferred over values farther away from a_1. The above procedure is applied to each non-zero security with a probability p_m.

Since the values are changed by this mutation operator, the repair mechanism (described in subsection 3.1) is applied to the mutated solution.

3.3.2 Mutation 2

The above mutation procedure does not convert a zero value to a non-zero value. But, by the second mutation operator, we change s securities to non-zero securities. To keep the number of investments d same, we also change s non-zero securities to zero securities. For this purpose, a zero and a non-zero security are chosen at random from a solution and their values are interchanged. The same procedure is re-applied to another $(s-1)$ pairs of zero and non-zero securities. To illustrate, we set $s = 2$ and apply the operator on the following solution:

$$a_1 \quad 0 \quad a_3 \quad 0 \quad a_5 \quad 0 \quad 0 \quad a_8 \quad 0 \quad a_{10}$$

Mutation 2 operator is applied between first and fourth securities and between ninth and tenth securities as follows:

$$0 \quad 0 \quad a_3 \quad a_1 \quad a_5 \quad 0 \quad 0 \quad a_8 \quad a_{10} \quad 0$$

Since no new non-zero value is created by this process, the 1st constraint will always be satisfied and this mutation operator always preserves feasibility of the solution.

Thus, the proposed customized hybrid NSGA-II procedure starts with a set of feasible solutions. Thereafter, the creation of new solutions by recombination and mutation operators always ensures their feasibility. These procedures help NSGA-II converge close to the Pareto front trading off the two conflicting objectives quickly.

3.4 Clustering

NSGA-II is expected to find a set of solutions representing the trade-off frontier. The NSGA-II procedure requires a population size (N) proportional to the size of the problem (total number of securities) for a proper working of its operators. However, this can cause NSGA-II to find a more portfolios than a user may desire. To reduce the number of obtained trade-off portfolios, we apply a k-mean clustering procedure to the obtained set of solutions and pick only H $(< N)$ well-distributed portfolios in risk-return space from the obtained NSGA-II solutions. For all our problems, we use $H = 30$ to get an idea of the trade-off frontier, however in practice a smaller number of solutions could be employed.

In addition to convergence to optimal solutions, a major emphasis of the NSGA-II procedure is to find a well-diversified set of points. Thus, a computationally efficient procedure would be to not use NSGA-II for a large number of generations, but rather pursue a hybrid strategy that improves each clustered NSGA-II solution by means of a local search procedure. We describe the local search procedure in the following.

3.5 Local Search

Due to the discontinuities in the constraints described earlier, we employ a single-objective GA for this purpose. From a seed solution $\mathbf{z}$ (a clustered solution obtained by NSGA-II), first an initial population is created by mutating the seed solution N_{ls} times (where N_{ls} is the population size of the local search procedure). Both mutation operators discussed above are used for this purpose. The mutation probability p_m and index η_m are kept the same as before. Since the mutated solutions are repaired, each population member is guaranteed to be feasible.

The objective function used here is the following achievement scalarizing function [24]:

$$F(\mathbf{x}) = \min_{j=1}^{2} \frac{f_j(\mathbf{x}) - z_j}{f_j^{\max} - f_j^{\min}} + 10^{-4} \sum_{j=1}^{2} \frac{f_j(\mathbf{x}) - z_j}{f_j^{\max} - f_j^{\min}}. \tag{8}$$

Here, $f_2(\mathbf{x})$ is considered to be the negative of the return function (second objective) in (3). The parameters $f_j^{\min}$ and $f_j^{\max}$ are the minimum and maximum objective values of the obtained NSGA-II solutions. In the local search GA, the above function is minimized.

Identical crossover and mutation operators (as described above) are applied to create a new population. The search is continued until no significant improvement in consecutive population-best solutions is obtained.

4 Results

We are now ready to present some results using our proposed customized hybrid NSGA-II procedure. In our experiments, we utilized a data-set (containing an $\mathbf{r}$ return vector and a $[\sigma_{ij}]$ covariance matrix) from a 1,000-security problem. The data-set was obtained from the random portfolio-selection problem generator specified in [13]. Figure 2 shows the risk-return front obtained by solving the ϵ-constraint formulation for the problem many times via QP. Solutions that

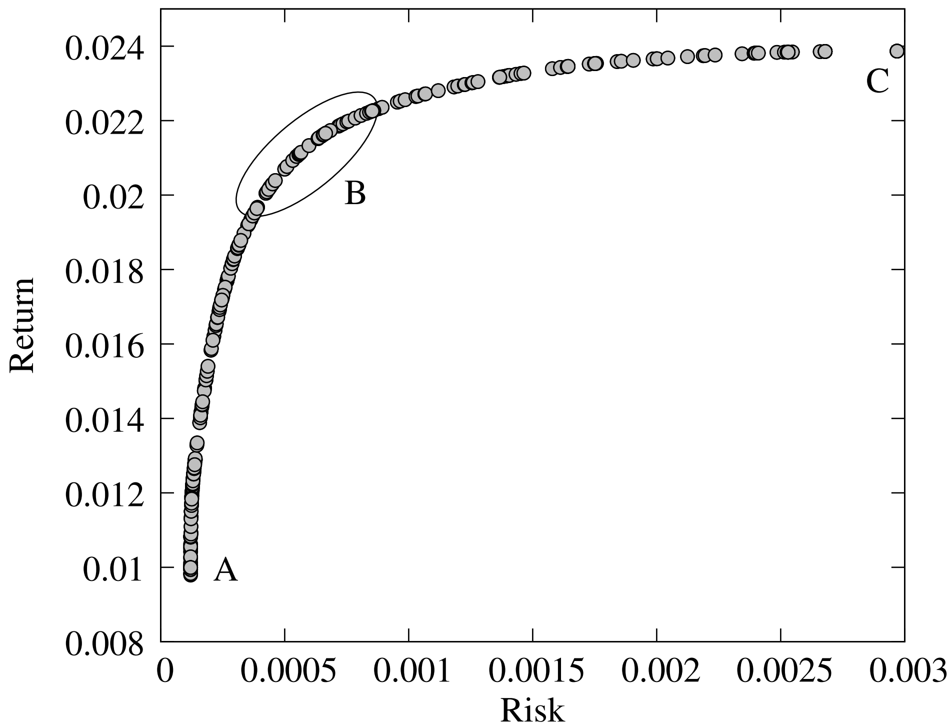

Figure 2: Risk-return front obtained by solving the ϵ-constraint formulation many times via QP

have low risk (near A) also have low expected returns. Solutions that have high risk (near C) also have a high expected returns. Consequently solutions near B are tempting but one can not tell for sure where a given investor's optimal solution might lie ahead of time. Thus, it is customary in portfolio selection to show the entire Pareto-front to give the user a global view of all potentially optimal solutions. In this way, a user is able to select his or her most preferred portfolio in full view of all other solutions knowing that nothing is hidden or had gotten lost in the process.

For this case study, we analyzed all ϵ-constraint QP solutions and observed that out of 1,000 securities in the problem, only 88 ever had non-zero investments. We extract these securities and their $\mathbf{r}$ and σ risk matrix information and confine our further studies to only these $n = 88$ securities.

For the NSGA-II procedure, we have chosen a population size of $N = 500$, a maximum number of generations of 3,000, a recombination probability of $p_c = 0.9$, an SBX index of $\eta_c = 10$, a mutation probability $p_m = 0.01$, and a polynomial mutation index of $\eta_m = 50$. For the mutation operator, we also set $s = 1$. For the local search GA, we used $N = 200$ and ran for a maximum of 500 generations. All other GA parameters are kept the same as in NSGA-II simulations.

4.1 Portfolio Optimization with No Size Constraint

Without any discontinuous constraints, the Pareto-front of a mean-variance problem is a connected collection of curved arcs each coming from a different parabola. Let us now consider the points at which one arc ends and another begins. Since each of these points has an R-value, in our experiments, for purposes of standardization, we let these be the R-values used in (2) to solve for the front in classical QP. Without a 3rd constraint (i.e., no size constraint), and $\alpha = 0$ and $\beta = 0.04$ in the 2nd, our 1,000-security problem has 483 such QP produced points. This is usually many more points than would ever be produced by QP, but this is good for the testing of our procedures. However, only 453 of them satisfy the 2nd constraint when $\alpha = 0.005$. Due to the use of $\beta = 0.04$, we note that at least 25 securities must exist in every feasible portfolio. However, in our algorithm, we do not specify any restriction on the number of securities.

Figure 3 shows 294 trade-off solutions obtained by the hybrid NSGA-II procedure alone. The QP solutions that satisfy the bound constraint are also shown in the figure. The figure shows that the proposed procedure is able to find a good set of solutions close to the points obtained by the QP procedure. High risk solutions are somewhat worse than that obtained by the QP method. But we shall show later that the hybrid procedure with a local search run from each of the NSGA-II solutions is a better overall strategy.

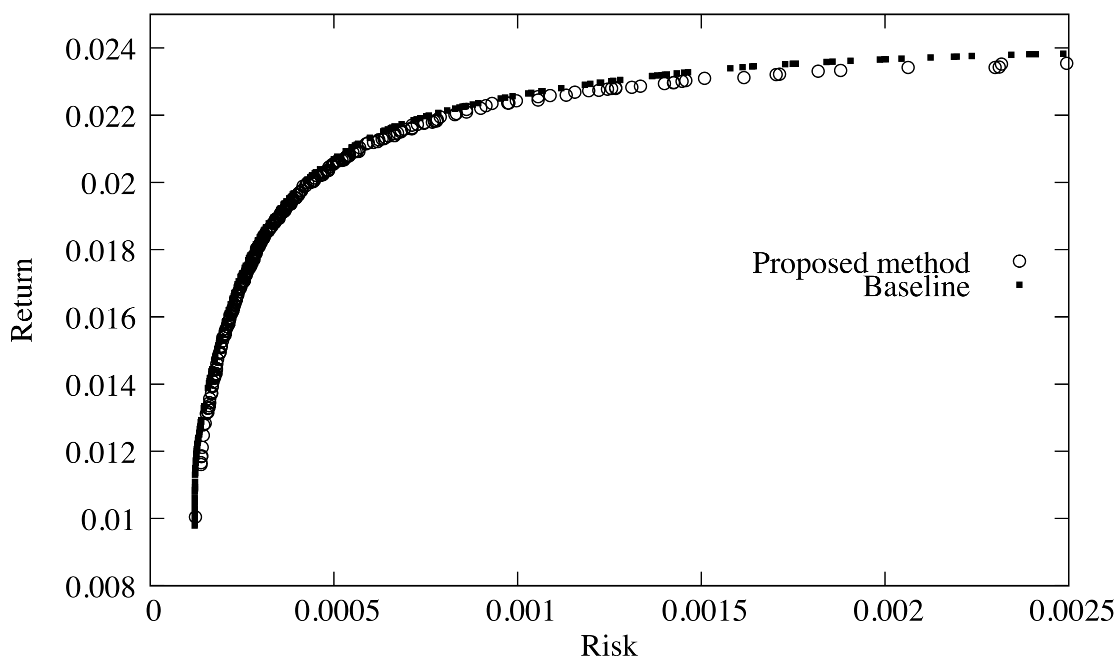

Figure 3: Obtained trade-off front for $\alpha = 0.005$ and $\beta = 0.04$, and with no restriction on d. Baseline means from the QP solution set

Genetic algorithms are stochastic optimization procedures and their performance is shown to depend somewhat on the initial population. In order to demonstrate the robustness of the proposed customized hybrid NSGA-II procedure over variations in the initial population, we created 21 random initial populations and ran the proposed algorithm independently from each one. The obtained NSGA-II solutions are collected together and 0%, 50% and 100% attainment curves [11] are plotted in Figure 4. The 0% attainment curve indicates the trade-off boundary which dominates all obtained trade-off solutions. The 100% attainment curve indicates the trade-off boundary that is dominated by all obtained trade-off solutions. The 50% attainment curve is that trade-off boundary that dominates 50% of the obtained trade-off solutions. The three curves obtained from the 21 non-dominated fronts are so close to each other that they cannot

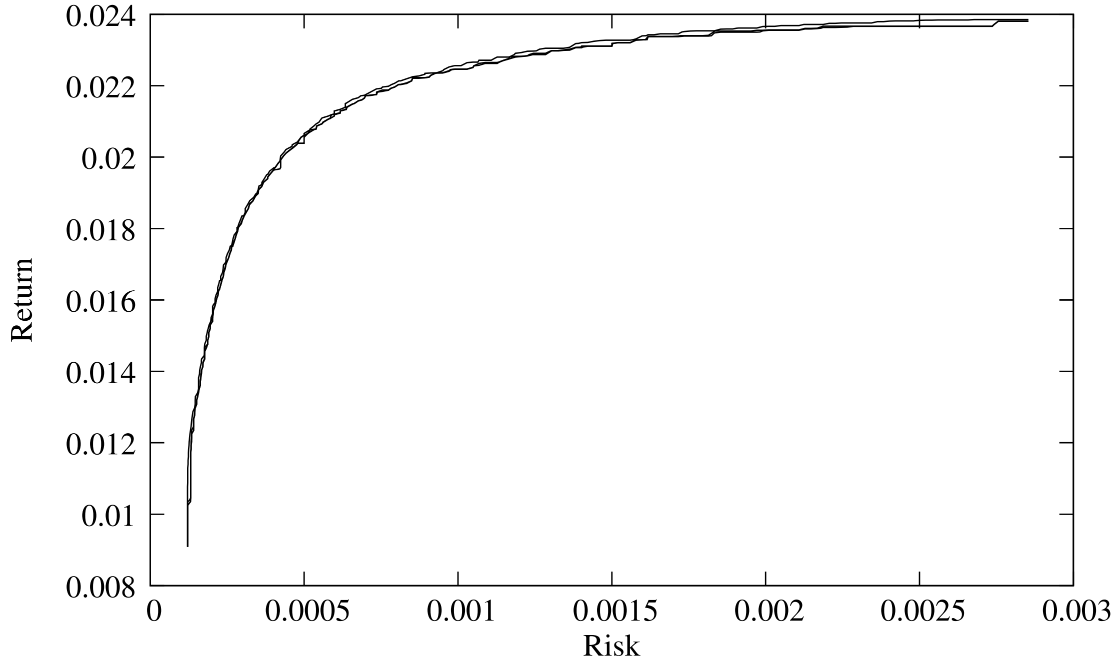

Figure 4: Attainment curves of the 21 runs for the 1,000-security problem for the case of $\alpha = 0.005$ and $\beta = 0.04$, and no restriction on d.

be well distinguished visually. This means that all 21 independent runs produce almost identical trade-off frontiers, thereby indicating the reliability of the proposed customized hybrid NSGA-II procedure.

From the obtained non-dominated front of all 21 runs, we now apply the clustering algorithm and choose only 30 well-distinguished portfolios covering the entire front. These solutions are then local searched and shown in Figure 5 by circles.

In order to show the investment pattern of different trade-off portfolios, we choose the two extreme solutions and two intermediate solutions from the obtained set of 30 portfolios. Based on their expected return values (r_i), all securities are ordered from the highest expected return value to the lowest. They are then divided into five grades of 18 securities each. The fifth grade has the lowest 16 securities. For each of the four solutions, the number of non-zero securities in each grade is counted and plotted in Figure 6. It can be clearly seen that the highest expected return solution allocates the maximum number of grade 1 securities in order to maximize the overall return. The solution with the lowest expected return (or lowest risk) invests more into grade 5 securities (having lowest expected return). The intermediate solutions invest more into intermediate grades to make a good compromise between return and risk.

We now consider a more difficult to achieve interval on acceptable proportions. Instead of $\alpha = 0.005$, we set $\alpha = 0.02$ and keep $\beta = 0.04$. This makes it more difficult for an optimization algorithm to find feasible solutions. Figure 7 shows the final set of 30 well-distributed, non-dominated points obtained by the NSGA-II, clustering and local search operations. Only five feasible solutions were found in the QP set. It is interesting to note that with a tighter bound on the proportions of investment, the number of invested securities is smaller. With an upper bound of 4% ($\beta = 0.04$), a minimum of 25 securities are to be invested. We observe that when $\alpha = 0.005$ was used, the average d for all 30 portfolios was 32.963. On the other hand, when $\alpha = 0.02$ is used, the average d is 27.034.

Finally, we repeat the study with $\alpha = 0.04$, identical to the upper bound. Not surprisingly, with $\alpha = \beta = 0.04$, all portfolios are found to have exactly 25 invested securities.

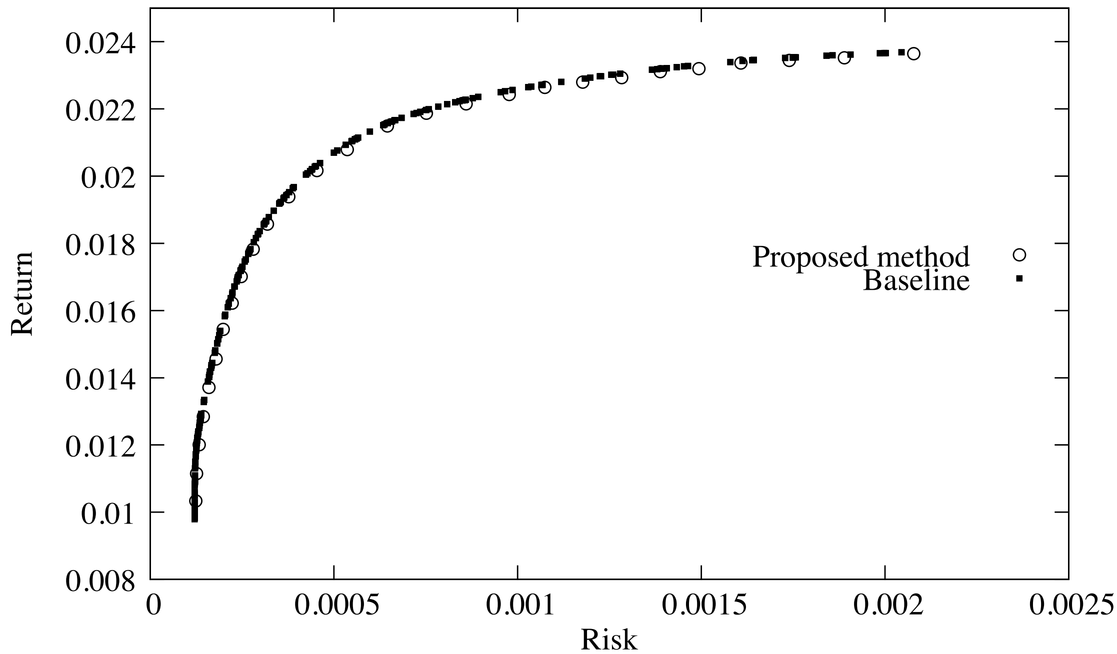

Figure 5: A set of 30 trade-off portfolios obtained using clustering and local search procedure for the case with $\alpha = 0.005$, $\beta = 0.04$, and no restriction on d.

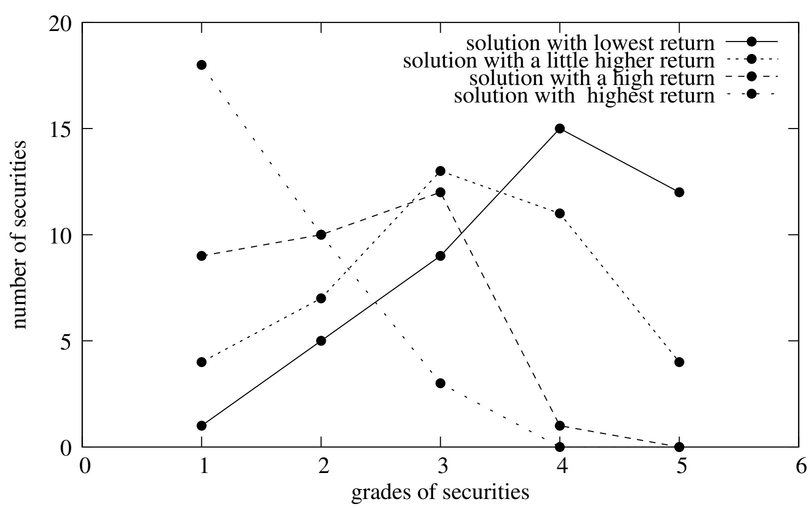

Figure 6: Investment pattern for four solutions in each of the five grades of diminishing expected return.

Figure 8 shows the obtained trade-off solutions for three different α values (0.005, 0.02 and 0.04), while the upper bound is kept the same, that is, at $\beta = 0.04$. The figure reveals an interesting feature of the risk-return portfolio optimization problem. Despite the tightness in the allowable bounds on the invested proportions, the ranges of risk and return values of the trade-off solutions are more or less the same. A large or a small amount of investment can be made to a small or a large number of securities, respectively, but this study shows that a variety of trade-off portfolios exists to span a similar range of associated risk and return values. Whether this is a feature of the particular securities (with their characteristic r_j and σ_{ij} values) considered here or of the particular upper bound ($\beta = 0.04$) considered here would be an interesting future study.

Before we close this section, we would like to reveal one more aspect of the risk-return portfolio

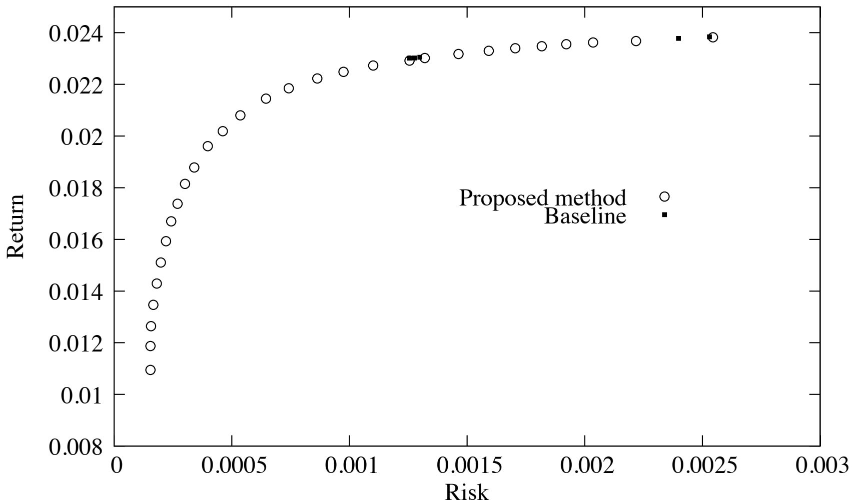

Figure 7: A set of 30 trade-off portfolios obtained using clustering and local search procedure for the case with $\alpha = 0.02$, $\beta = 0.04$, and no restriction on d.

optimization problem that most multi-objective optimization problems share with one another. The first author has conjectured and demonstrated on a number of bi-objective practical optimization problems that Pareto-optimal solutions tend to share some common characteristics. The so-called *innovization* task is a process by which a serial application of a multi-objective optimization algorithm finds a set of Pareto-optimal solutions and an analysis of the solutions unveils a set of common characteristics hidden in the solutions either manually [10] or automatically [1]. For the 30 wide-spread trade-off solutions obtained with $\alpha = \beta = 0.04$, we perform a manual analysis of the solutions to look for any such hidden characteristics or properties. We observe that out of the total 88 securities, 13 of them are not chosen for investment in any of the 30 optimized portfolios and one of them is chosen for investment in 29 of 30 portfolios. Figure 9 shows the number of optimized portfolios (of 30) in which each security is invested. The securities are arranged on the abscissa according to their expected return values. The left-most security (ID=0) corresponds to the lowest expected return security and the right-most security (ID=88) corresponds to the highest expected return security. The figure also plots the cumulative risk value for i-th security ($\sum_j \sigma_{ij}$) to get an idea of the overall risk associated with invested securities. Several principles can be derived from this figure:

1. Securities with higher expected return values are chosen more often in the optimized portfolios. All securities invested in more than 50% of the 30 optimized portfolios are associated with high expected return values.

2. Similarly, securities with smaller expected return values are not chosen for investment often. Of the 13 non-invested securities (in any of the 30 optimized portfolios), 11 of them are associated with low expected return values.

3. Not all high expected return securities are chosen equally. Of the high expected return securities, the ones with relatively high risk are invested in sparingly. High risk securities such as the ones corresponding to A and B, are not often chosen.

4. Even low or medium expected return but relatively high risk securities (C, D and so on) are not chosen often. In general, low risk securities are chosen more often.

These common properties of the trade-off optimized portfolios provide useful information to an user. Working out such an innovization analysis for a number of other scenarios should help

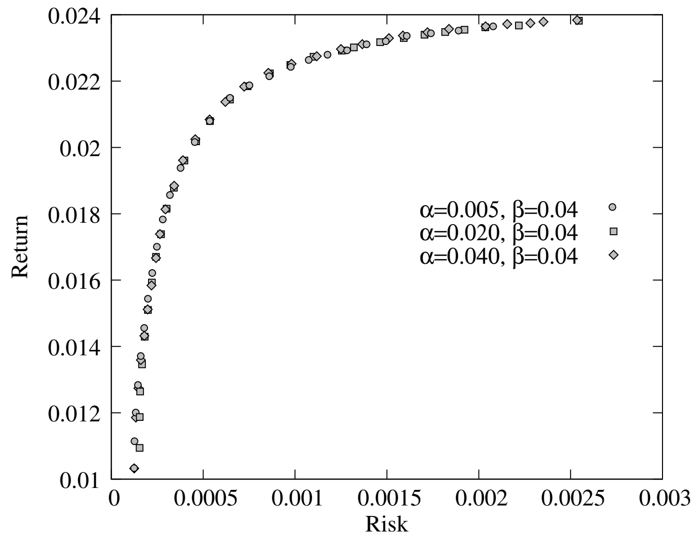

Figure 8: The effect of lower bound (α) on the trade-off portfolios indicates no substantial change in the range of risk-return values.

generate a set of 'thumb rules' for performing an efficient portfolio management task.

4.2 Portfolio Optimization for a Fixed d

Now we focus on the 3rd constraint. Initial population members are guaranteed to satisfy the 3rd constraint and all recombination and mutation operations also guarantee that the 3rd constraint is satisfied.

First, we consider $d^{\min} = d^{\max} = 28$, so that portfolios with only 28 securities are desired. The variable bounds of $\alpha = 0.005$ and $\beta = 0.04$ are enforced. The QP generated frontier now has only 24 portfolios which satisfy these constraints. They are shown in Figure 10. Obtained NSGA-II solutions are also marked on the figure. There are a total of 168 portfolios found by a typical run of the procedure. It is clear from the figure that the proposed method is able to find a widely distributed set of solutions.

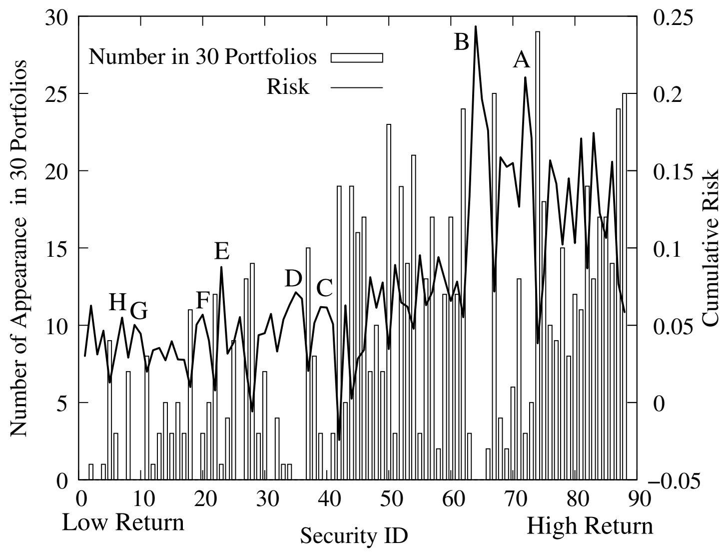

Figure 9: An innovization task for deciphering hidden investment characteristics or properties common to the 30 optimized portfolios.

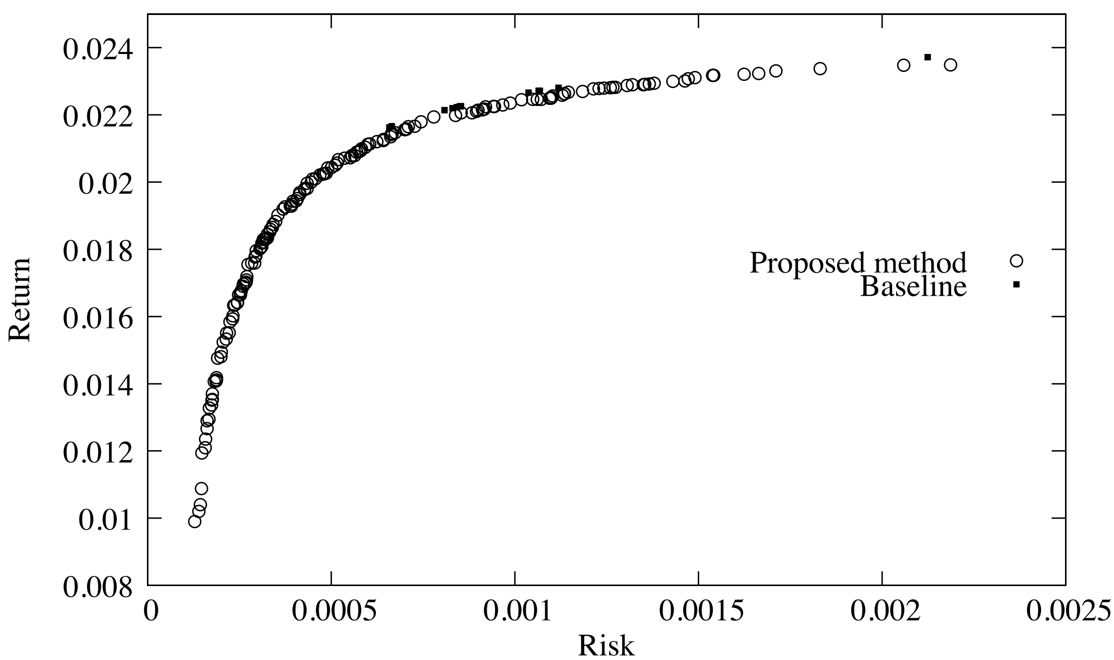

Figure 10: Obtained trade-off front for $d = 28$, with $\alpha = 0.005$ and $\beta = 0.04$.

In order to investigate the robustness of the NSGA-II procedure alone, the three attainment curves are drawn for 21 different runs in Figure 11, as before. The closeness of these curves to one another signifies that the proposed procedure is able to find almost identical trade-off frontiers in each run, thereby demonstrating the robustness of the proposed procedure.

Finally, all solutions from the 21 runs are collected and the non-dominated solutions are clustered into 30 widely separated groups. After being local searched, we have the results shown in Figure 12. It is clear from the figure that the proposed methodology is able to find evenly dispersed sets of portfolios.

In order to understand how the NSGA-II driven procedure is able to find feasible solutions

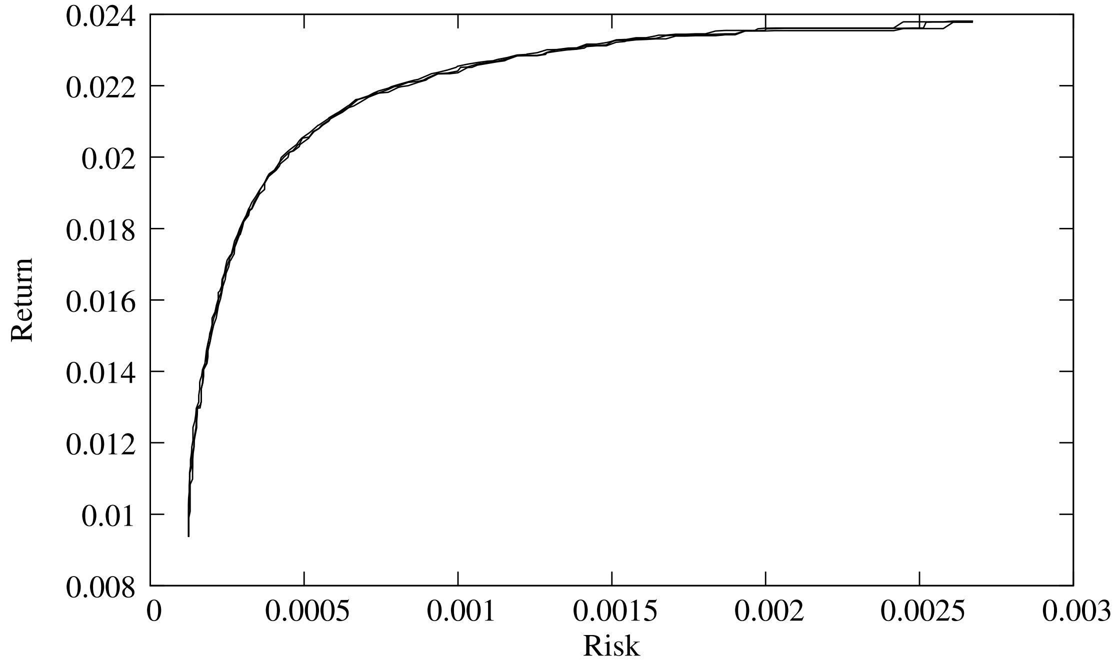

Figure 11: Attainment curves for the 21 runs with $d = 28$, and $\alpha = 0.005$ and $\beta = 0.04$.

satisfying the constraints, we pick two solutions from the NSGA-II local-search front and find the two solutions in risk-return space closest to the QP front. The corresponding solutions are shown in Figures 13 and 14. In the first figure, the corresponding objective values are as follows: NSGA-II: risk=0.000176, return=0.014352; QP: risk=0.000176, return=0.014743. Although they are close in their objective values, the QP solution is infeasible on two counts. First, the number of invested securities is not equal to the desired number (28). Rather, the portfolio has 35. Second, the QP solution has one investment with a proportion of 0.004003, which of course, is not allowed, due to its violating the allowable lower bound of $\alpha = 0.005$.

On the other hand, the NSGA-II solution has exactly 28 invested securities and no investments in the gap between 0 and α. To understand how the NSGA-II solution is different from this QP solution, we plot the investment pattern of each security in an interesting manner in Figure 13. Every point in the figure corresponds to the respective proportion of investment for a particular security in the QP portfolio (abscissa) and in the NSGA-II portfolio (ordinate). The NSGA-II solution converts some of the QP invested securities to non-investment status (like securities A), while some of the non-invested securities are converted to investment status (like securities B and C). Also, some intermediate QP investments are converted to maximum allowable investment status (securities D), and some maximum allowable investments are converted to intermediate status (like securities E). A few securities are changed from one intermediate value to another (like securities F). Statistics describing the changes in investment status are shown in Table 1.

It is interesting to note that a majority (about 57%) of the zero investment securities are preserved between the QP and corresponding NSGA-II solutions. Only three zero investment securities are converted to non-zero investment status, but 10 non-zero investment securities are converted to zero investment status. This is done so as not to deviate from the mandated 28. Also, a number of securities (14) with maximum allowable investment are reassigned the same way. The NSGA-II's customized operators make these changes in a way so that no investments occur in the range $(0, 0.005)$ and that there are exactly the requisite number of invested securities; yet the corresponding risk and return values are closer to those of the QP solution. Table 1 tabulates statistics on these adjustments.

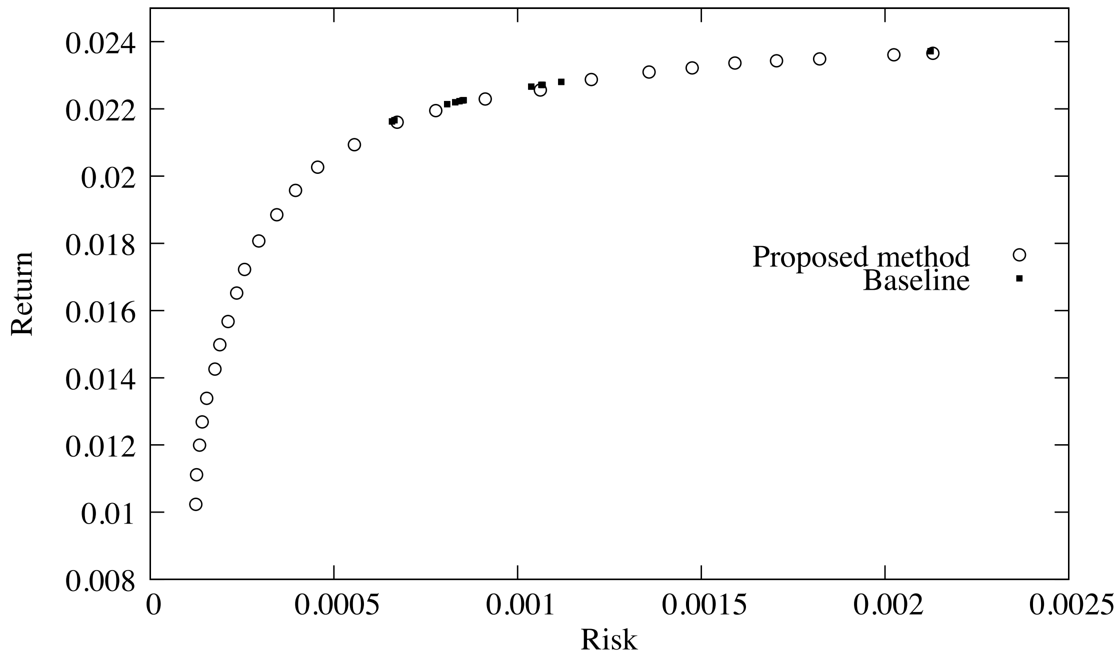

Figure 12: A set of 30 trade-off portfolios obtained using clustering and local search procedure for the case with $d = 28$, $\alpha = 0.005$, and $\beta = 0.04$.

Table 1: Difference between the infeasible (QP) and feasible (NSGA-II) solutions having similar objective values. Solution 1 has 13 securities, and Solution 2 has 6 securities, that change from investment to non-investment status. Whereas Solution 1 by QP has seven more securities than allowed (28), Solution 2 by QP has two more.

	Identical x_i		Different x_i		
	$x_i^{\text{QP}} = 0$ and $x_i^{\text{NSGA-II}} = 0$	$x_i^{\text{QP}} = 0.04$ and $x_i^{\text{NSGA-II}} = 0.04$	$x_i^{\text{QP}} = 0$ and $x_i^{\text{NSGA-II}} > 0$	$x_i^{\text{QP}} > 0$ and $x_i^{\text{NSGA-II}} = 0$	Otherwise
Sol. 1	50	3	10	14	11
Sol. 2	56	2	4	17	9

Another such example is shown in Figure 14, in which the risk and return values are close and as follows: NSGA-II: risk=0.000556, return=0.020932; QP: risk=0.000557, return=0.021092. The QP portfolio has 30 invested securities (against 28 allowed) and three invalid investments with proportions 0.000474, 0.003230, and 0.000645. On the contrary, the corresponding NSGA-II solution does not violate any of these constraints. The figure and Table 1 show the adjustments made to transform the QP solution (with two additional and three invalid investments) to a feasible solution by the proposed NSGA-II's customized operators.

To conclude this section, we consider another case in which exactly $d = 50$ securities are to be in each portfolio. We take this scenario, because the original QP solution set did not include any such portfolio. We investigate the proposed algorithm's ability to handle such a portfolio. When we apply our proposed clustered NSGA-II-local-search approach, we obtain the front with 30 portfolios shown in Figure 15. A wide spread of portfolios in terms of risk and return is found.

These results amply demonstrate the working of the proposed optimization procedure when both discontinuous variable bounds and a discrete number of securities are desired.

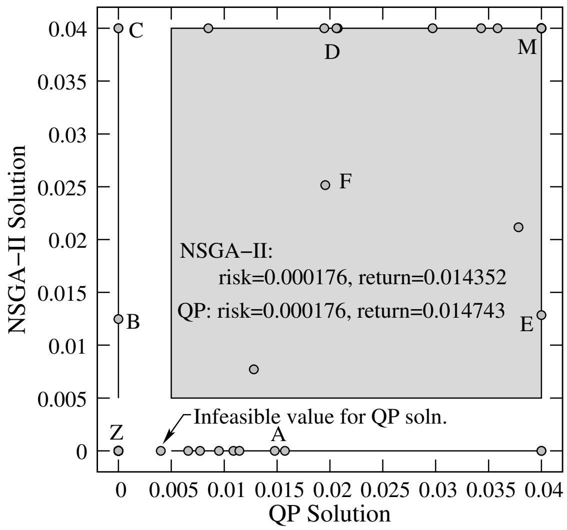

Figure 13: Two similar low-return portfolios – an infeasible one obtained by QP and another feasible obtained by the NSGA-II procedure – are illustrated.

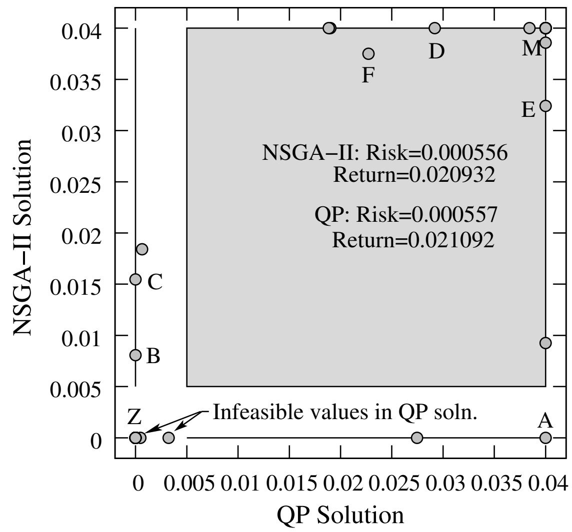

Figure 14: Two similar medium-return portfolios – an infeasible one obtained by QP and another feasible obtained by the NSGA-II procedure – are illustrated.

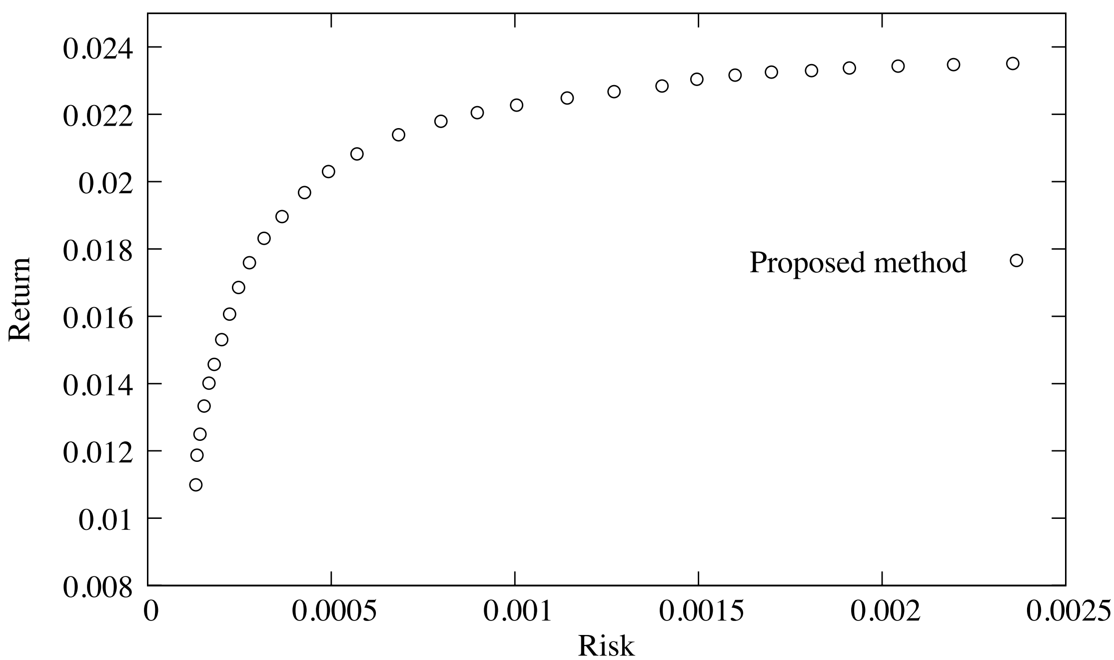

Figure 15: 30 different portfolios found for $d = 50$, $\alpha = 0.005$, and $\beta = 0.04$ obtained using the proposed approach. The QP solution set did not have any solutions with $d = 50$.

4.3 Portfolio Optimization for a Given Range of d

Now let us consider the case in which the user wishes to have portfolios with investments in the prescribed range $[d^{\min}, d^{\max}]$, instead of investing exactly in a specified number securities. The NSGA-II procedure to handle such a scenario is described before.

Let us consider $d^{\min} = 30$ and $d^{\max} = 45$ and select the variable bounds to be $\alpha = 0.005$ and $\beta = 0.04$. In this problem, first we show the effect of the local search. Figure 16 shows

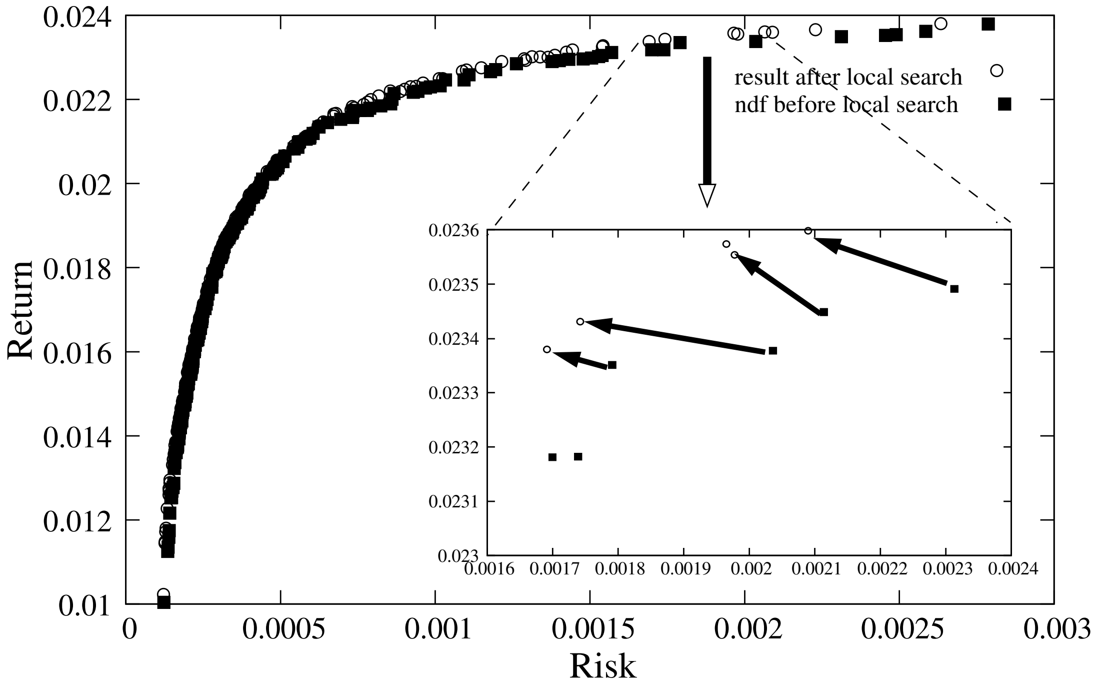

Figure 16: Effect of local search for $\alpha = 0.005$, $\beta = 0.04$, and $d \in [30, 45]$.

the solutions before and after the local search on a set of trade-off solutions obtained in a single NSGA-II simulation. One can see the improvements made by local search in the figure.

To demonstrate the robustness of the proposed approach, we now perform 21 different NSGA-II simulations. As shown in Figure 17, the three attainment derived from the 21 runs are very close to one another, thereby indicating the all 21 runs find almost identical trade-off frontiers.

Finally, all solutions from the 21 runs are collected and the non-dominated solutions are clustered into 30 widely separated groups. Then after being local searched, we have the results shown in Figure 18. In the range for d, 261 of the QP obtained solutions are found to be feasible. These solutions are marked in the figure. The closeness of the computed solutions to the QP solutions is clear in the figure.

In order to understand the pattern of investment vis-a-vis the number of invested securities (d), we plot d for each solution along with its corresponding risk value (objective $f_1(\mathbf{x})$) in Figure 19. It is interesting to note that low risk portfolios (also with low returns) can be obtained with a variety of investments, almost uniformly from 30 to 44 specific securities. However, high return portfolios (also with high risk) require not more than a small number of securities.

Next, we analyze the portfolios and attempt to understand why the high and low risk portfolios are made up of the securities that they are. For a specific high risk portfolio and a specific low risk portfolio from the final trade-off frontier, we plot the proportions of investment made to the different securities in Figure 20. First, the securities are ordered from high return (abscissa equal to 1) to low return (abscissa equal to 88). It is clear that the high risk portfolios invest more in high-return securities and the low risk portfolios invest more in low return securities. Although this outcome is intuitive, it is interesting that the proposed optimization algorithm is able to bring out this feature of investment, thereby providing confidence in the working behavior of the proposed procedure.

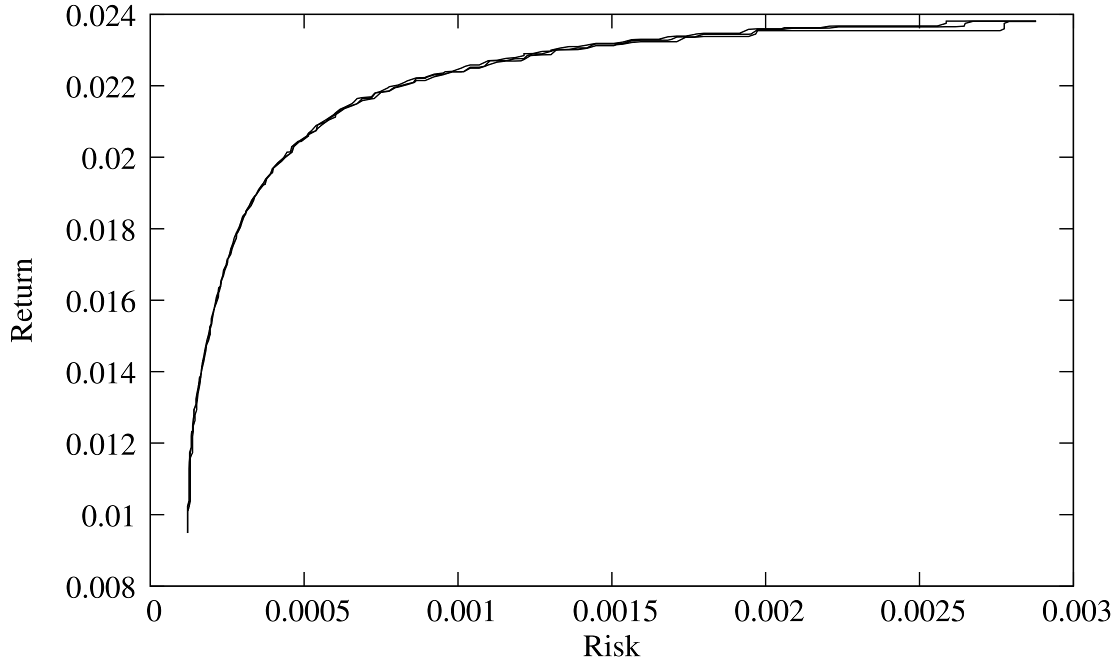

Figure 17: Attainment curves for the 21 runs when $d \in [30, 45]$, $\alpha = 0.005$, and $\beta = 0.04$.

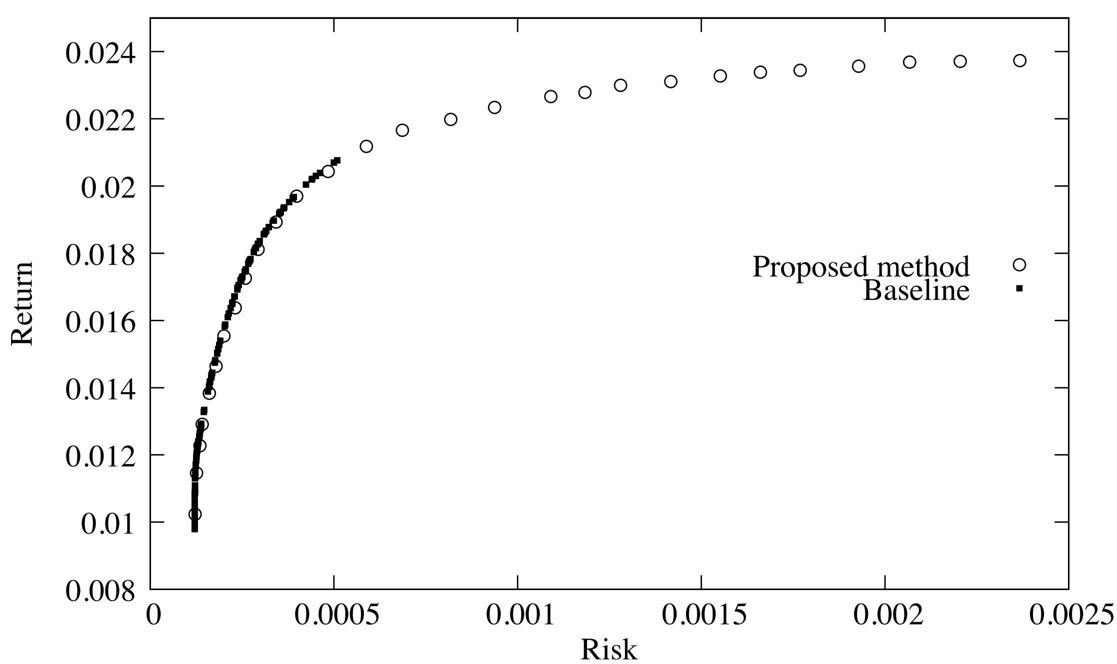

Figure 18: A set of 30 portfolios obtained for $d \in [30, 45]$, $\alpha = 0.005$, and $\beta = 0.04$.

5 Conclusions

In this paper, we have suggested a customized hybrid NSGA-II procedure for handling a bi-objective portfolio optimization problem having practicalities. First, the proportion of investment in a security can either be zero (meaning no investment) or be between specified minimum and maximum bounds. Second, the user may have a preference for a certain number or range of invested securities. Such constraints make the resulting optimization problem difficult to solve

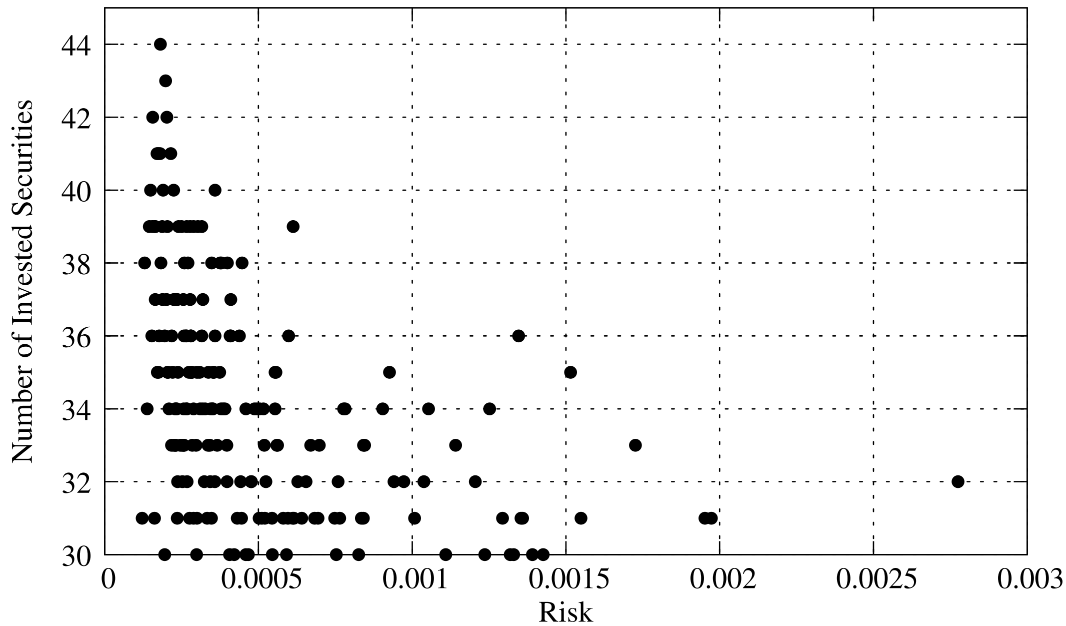

Figure 19: Number of invested securities (d) versus associated risk in 30 optimized portfolios. For high risk portfolios, options are limited. For low risk portfolios, options are many.

using classical QP methods. The hallmark of our study is that we have suggested a customized procedure that is able to repair solutions so as to make them feasible.

By taking a large-sized problem, we have systematically demonstrated the efficacy of the proposed procedure over different user-defined requirements. The robustness of the proposed customized hybrid NSGA-II is demonstrated by simulating the algorithm from 21 different initial populations. In all cases, the attainment curve plots indicate that all 21 runs find virtually identical trade-off frontiers. The accuracy of the proposed procedure is enhanced by coupling the method with clustering and local search components. The results are compared with those from a QP method and the following conclusions can be drawn:

- The proposed customized hybrid NSGA-II and local search procedure is able to find a wide variety of feasible portfolios honoring different desired practicalities.
- Although a feasible solution may exist close to an infeasible QP solution, the difference in the investment pattern between the solutions can be substantial. This may require a number of previously non-invested securities to become invested, and a number of previously invested securities to become non-invested. Also, changes in the amounts invested in the securities that remain invested may be necessary.
- The range on allowed investments dictates the number of optimized securities that must be invested. For a tighter range for investment, investments must be made on fewer securities, and vice versa.

Besides being able to handle practicalities associated with the portfolio optimization problem, the bi-objective trade-off portfolios are analyzed to decipher important properties common to them:

- High return (and high risk) optimized portfolios are obtained by investing more in high return and relatively high risk securities.

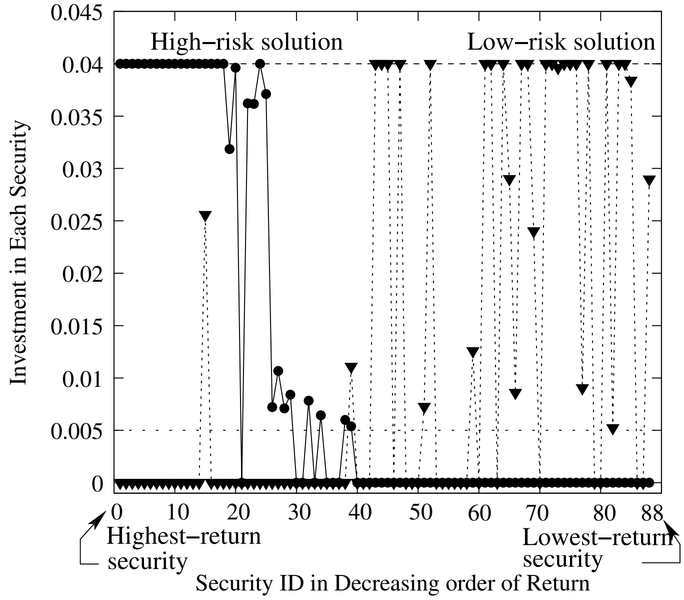

Figure 20: Proportions of investments in each security for a high and a low risk portfolio.

- Low return (and low risk) optimized portfolios are obtained by investing more on low return and relatively low risk securities.
- To obtain a high return optimized portfolio, in general, investments to a few high-return securities are needed.
- On the other hand, to obtain low-return portfolios, in general, we are relegated to only a few securities.
- Interestingly, a significant number of securities are found to be commonly non-invested in the entire set of trade-off portfolios. These securities have varying degrees of return values, but one common aspect of these securities is that they all carry relatively higher risk.

The methodology suggested here is now ready to be applied to more complex portfolio optimization problems. The QP method, when applicable, is unmatched by any other algorithm. But practice is full of non-linearities and complexities which prohibit the use of standard QP methods alone. The flexibility of genetic algorithms as demonstrated in this paper provides us with a confidence in their ability to deal with more real-life portfolio optimization problems. For the future, a more pragmatic and computationally efficient approach would be to combine QP and GA by somehow taking advantage of the comparative advantage of each.

Acknowledgements

This work is funded by the Academy of Finland under project number 133387. Authors appreciate the support.

References

[1] S. Bandaru and K. Deb. Towards automating the discovery of certain innovative design principles through a clustering based optimization technique. *Engineering optimization*, in press.

[2] J. Branke, B. Scheckenbach, M. Stein, K. Deb, and H. Schmeck. Portfolio optimization and an envelope-based multi-objective evolutionary algorithm. *European Journal of Operational Research*, 199(3):684–693, 2009.

[3] T.-J. Chang, N. Meade, J. B. Beasley, and Y. Sharaiha. Heuristics for cardinality constrained portfolio optimisation. *Computers & Operations Research*, 27:1271–1302, 2000.

[4] Y. Crama and M. Schyns. Simulated annealing for complex portfolio selection problems. *European Journal of Operational Research*, 150:546–571, 2003.

[5] K. Deb. An efficient constraint handling method for genetic algorithms. *Computer Methods in Applied Mechanics and Engineering*, 186(2–4):311–338, 2000.

[6] K. Deb. *Multi-objective optimization using evolutionary algorithms.* Wiley, Chichester, UK, 2001.

[7] K. Deb and R. B. Agrawal. Simulated binary crossover for continuous search space. *Complex Systems*, 9(2):115–148, 1995.

[8] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast and elitist multi-objective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

[9] K. Deb and A. Kumar. Real-coded genetic algorithms with simulated binary crossover: Studies on multi-modal and multi-objective problems. *Complex Systems*, 9(6):431–454, 1995.

[10] K. Deb and A. Srinivasan. Innovization: Innovating design principles through optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2006)*, pages 1629–1636, New York: ACM, 2006.

[11] C. M. Fonseca, V. Grunert da Fonseca, and L. Paquete. Exploring the performance of stochastic multiobjective optimisers with the second-order attainment function. In *Third International Conference on Evolutionary Multi-Criterion Optimization, EMO-2005*, pages 250–264. Berlin: Springer, 2005.

[12] D. E. Goldberg. *Genetic Algorithms for Search, Optimization, and Machine Learning.* Reading, MA: Addison-Wesley, 1989.

[13] M. Hirschberger, Y. Qi, and R. E. Steuer. Randomly generating portfolio-selection covariance matrices with specified distributional characteristics. *European Journal of Operational Research*, 177(3):1610–1625, 2007.

[14] T. Krink and S. Paterlini. Differential evolution for multiobjective portfolio optimization. Center for Economic Research (RECent) 021, University of Modena and Reggio E., Dept. of Economics, 2008.

[15] C.-M. Lin and M. Gen. An effective decision-based genetic algorithm approach to multiobjective portfolio optimization problem. *Applied Mathematical Sciences*, 1(5):201–210, 2007.

[16] D. Lin, X. Li, and M. Li. A genetic algorithm for solving portfolio optimization problems with transaction costs and minimum transaction lots. In *Lecture Notes in Computer Science, 3612*, pages 808–811, Berlin: Springer, 2005.

[17] H. M. Markowitz. *Portfolio selection: Efficient diversification of investments.* Yale University Press, New York, 1959.

[18] M. Stein, J. Branke, and H. Schmeck. Efficient implementation of an active set algorithm for large-scale portfolio selection. *Computers & Operations Research*, 35(12):3945–3961, 2008.

[19] Michael Stein, Jrgen Branke, and Hartmut Schmeck. Portfolio selection: How to integrate complex constraints. *Journal of Financial Planning*, pages 68–78, 2005.

[20] R. E. Steuer. *Multiple Criteria Optimization: Theory, Computation and Application.* New York: Wiley, 1986.

[21] R. E. Steuer, Y. Qi, and M. Hirschberger. Suitable-portfolio investors, nondominated frontier sensitivity, and the effect of multiple objectives on standard portfolio selection. *Annals of Operations Research*, 152(1):297–317, 2007.

[22] F. Streichert and M. Tanaka-Yamawaki. The effect of local search on the constrained portfolio selection problem. In *IEEE International Congress on Evolutionary Computing (CEC 2006)*, pages 2368–2374. Piscatway: IEEE Press, 2006.

[23] F. Streichert and M. Tanaka-Yamawaki. A memetic algorithm on the constrained portfolio selection problem. *Institute of Statistical Mathematics*, 187:140–149, 2006.

[24] A. P. Wierzbicki. The use of reference objectives in multiobjective optimization. In G. Fandel and T. Gal, editors, *Multiple Criteria Decision Making Theory and Applications*, pages 468–486. Berlin: Springer-Verlag, 1980.

[25] W.-G. Zhang, W. Chen, and Y.-L. Wang. The adaptive genetic algorithms for portfolio selection problem. *IJCSNS International Journal of Computer Science and Network Security*, 6(1):196–200, 2006.