AN EVOLUTIONARY ALGORITHM FOR PORTFOLIO SELECTION IN A DOWNSIDE RISK FRAMEWORK

Andrea Loraschi
SIGE Consulenza SIM S.p.A., C.so Matteotti 4/6 — 20124 Milano, ITALY
Tel.: +39 2 7751 2608
E-mail: A.loraschi@relay1.iunet.it

Andrea Tettamanzi
Università degli Studi di Milano, Dip. di Scienze dell'Informazione
Via Comelico 39/41 — 20135 Milano, ITALY
Tel.: +39 2 5500 6293
E-mail: tettaman@dsi.unimi.it

Abstract

An application of Evolutionary Algorithms to portfolio optimization is presented, which makes possible the use of nonstandard indices of risk. After a critical overview of Modern Portfolio Theory, portfolio optimization is formulated as a two-objective optimization problem and an evolutionary algorithm to solve it is described in some detail. A brief discussion of results is enclosed.

1. Introduction

This paper examines a financial application of Evolutionary Algorithms (Michalewicz 1992) to some issues raised by recent research in Portfolio Theory (PT) concerning the use of non-trivial indices of risk.
The central problem of PT concerns selecting the weights of assets in a portfolio which minimize a certain measure of risk for any given level of expected return.
According to mainstream PT (Markowitz 1959), a meaningful measure of risk is given by the variance of the distribution of portfolio returns. However, while leading to elegant analytical results and easy practical implementation, this approach fails to capture what an investor perceives as the essence of risk, that is the chance of incurring a loss.
In this respect the economic literature has recently proposed alternative measures of risk; an interesting family of measures, defined as lower partial moments (Harlow 1991), refers to the down-side part of the distribution of returns and has therefore become known under the name of downside risk.
A number of difficulties are encountered when trying to apply this new approach even to simple problem instances. First of all, the nature of these risk measures prevents from devising general analytical solutions and quadratic optimization techniques cannot be applied because the shape of the objective function is in general non-convex. Moreover, the typical sizes of real-world portfolio selection problems are of the order of hundreds of assets; expected return and risk are calculated using historical series made up of hundreds of historical returns. Even the software packages that apply the conventional mean-variance approach through quadratic optimization suffer limitations as the size of the problem grows beyond a certain threshold. Things become more involved as it is usual practice among fund managers to impose various constraints on their optimal portfolios.

The authors have carried out experiments applying Evolutionary Algorithms to the portfolio selection problem using the downside risk approach, obtaining satisfactory results. Evolutionary Algorithms appear to be well suited to the task of solving difficult problems in high dimensional spaces because they combine features from random search and Monte Carlo methods with some powerful heuristics borrowed from natural evolution. On the basis of this approach, the authors have developed a software decision support system for SIGE Consulenza SIM which has been named DRAGO (Downside Risk Asset allocation with Genetic Optimization). This tool has proved successful in devising portfolios where the optimized set of weights provides protection against potential underperformance of the portfolio with respect to a given benchmark.
Section 2 provides a critical overview of Modern Portfolio Theory and discusses a family of downside risk indices; Section 3 is a brief introduction to Evolutionary Algorithms; portfolio optimization is formulated as a two-objective convex optimization problem in Section 4 and an evolutionary algorithm to solve it (DRAGO) is described in Section 5; a brief discussion of experiments and result is given in Section 7.

2. Portfolio optimization

2.1. Modern Portfolio Theory

Choosing an optimal portfolio weighting of assets when their future rate of return is uncertain may be seen as a problem of minimizing the uncertainty for a given level of the portfolio expected return. Modern Portfolio Theory (MPT) calls this uncertainty risk and measures it as the standard deviation or the variance of the probability distribution of future returns. In this framework portfolios are optimized through the use of quadratic programming methods.
Other more recent approaches that require less simplistic assumptions, such as those based on the idea of "downside-risk", are not very easy to deal with using deterministic combinatorial optimization methods.
The notion of risk as the variance of returns is at the foundation of MPT. Yet, Harry Markowitz (1959), acknowledged as one of the first developers of MPT,
pointed out that variance

> *…often leads to unsatisfactory predictions of [investor] behavior.*

Realizing that the investor perceives risk only on the downside part of the return distribution, he proposed the use of "semivariance", defined as the weighted average of the squared distances of the returns *below* the mean of the distribution, as an index of risk. In mathematical notation:

$$\text{Semivar} = \sum_{R_i < E[R]} p_i (R_i - E[R])^2$$

where R_i is one of the elements of the portfolio return distribution, p_i its attached probability and $E[R]$ the expected (mean) return. However, because of the inherent difficulties in carrying out the necessary calculations, he built his model on variance. Choosing an optimal portfolio thus became a question of variance minimization for a given level of expected return. Markowitz showed that as the level of risk increases, the expected return attached to optimal portfolios draws a convex non-decreasing curve

(*efficient frontier*). That is, to attain a greater expected return, an investor must accept a higher exposure to risk.
Minimizing variance is equivalent to minimizing semivariance *if the return distribution is symmetric around the mean*. Distributions of returns, though, are more often skewed than symmetric: if, for example, a distribution is highly skewed to the left, variance will greatly underestimate the risk of realizing a return below the mean.

2.2. Drawbacks of Modern Portfolio Theory

Variance has the advantage of lending itself naturally to analytical treatment, but it shows serious drawbacks when used as an index of risk. As Domar and Musgrave (1944) remarked in a seminal article:

> *Of all the possible questions which the investor may ask, the most important one, it appears to us, is concerned with the probability of a loss. This is the essence of risk.*

Therefore, semivariance might be regarded as an alternative candidate, but its reliance on the concept of distance from the mean—since it varies from distribution to distribution—may lead to the unwanted result of discarding an asset which should be preferred to others because it could offer *with certainty* higher realized returns even if it shows a relatively higher semivariance. For instance, Fig. 1 shows two return distributions with same mean, but different variances: intuitively, the higher variance distribution will be, in fact, regarded as less risky by an investor, contrary to what MPT would suggest.

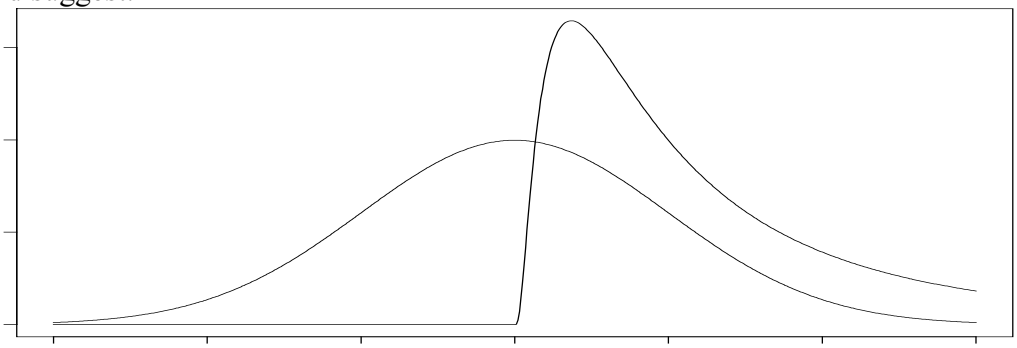

Fig. 1

2.3. Downside-risk indices

A new approach, which only recently has been partially explored by the literature in the field, defines risk as a measure related to the part of a return distribution below a certain target of return. Adapting the formula of semivariance to the concept of "subjective threshold below which the investor judges risky the investment" τ leads to the *target semivariance* expression:

$$\text{T.S.} = \sum_{R_i < \tau} p_i (\tau - R_i)^2 .$$

This measure of *downside risk* (Bawa 1975, Harlow1991) is usefully generalized as

$$\text{DSR}(\tau) = \sum_{R_i < \tau} p_i (\tau - R_i)^q .$$

The parameter q may be seen as a coefficient of risk aversion in the sense that if

- q=0, the above expression equals the probability of realizing a return below the target (*shortfall probability*);
- q=1, it shows the expected value of the returns below the target (*target shortfall*);
- q=2, it is equivalent to the *target semivariance*;
- q>2, the formula attaches increasing weight to the returns below the target, i.e. the higher the parameter, the more the investor dislikes downside-risk.

2.4. Implementation issues

A straightforward substitution of DSR in place of variance within the MPT context raises computational problems: this formula does not lend itself to analytical treatment and the classic deterministic algorithms used for optimization purposes are no longer suitable.
Fig. 2 shows the downside-risk surface for a simple problem involving only three assets: as it can be seen, the surface is very rugged and there are many local optima. A gradient descent method could not handle such a situation in a satisfactory way; moreover, as the number of dimensions increases, the computation effort required for calculating the gradient increases accordingly, and usual dimensions are in the order of the hundreds.

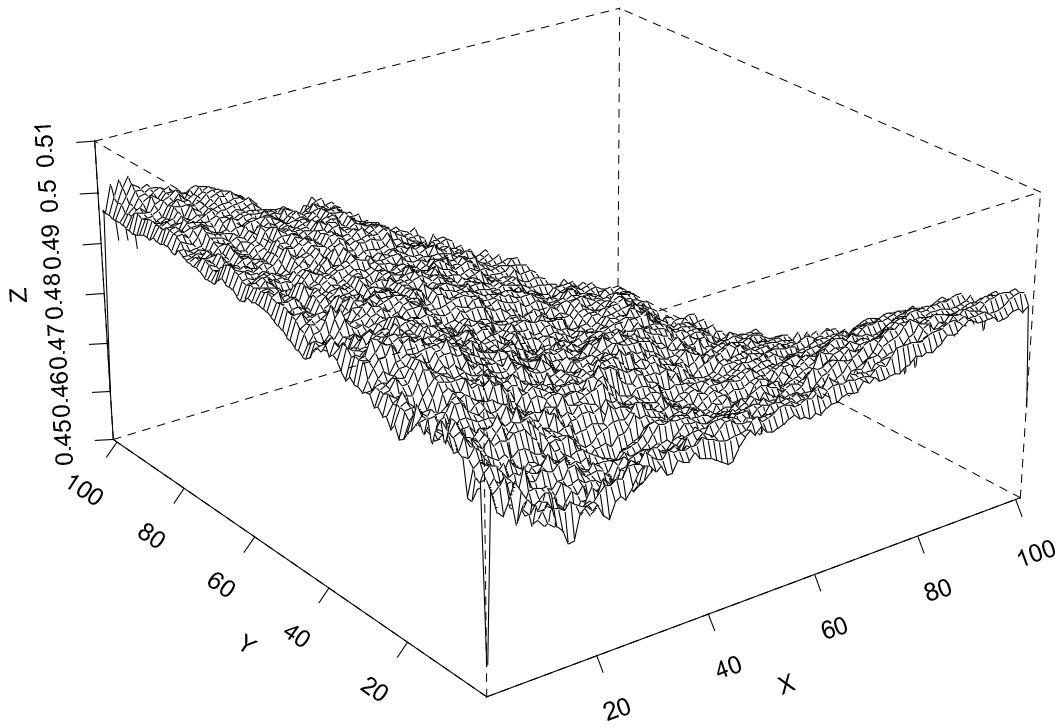

Fig. 2

3. Evolutionary Algorithms

The name *Evolutionary Algorithms* encompasses a family of stochastic optimization techniques based on the key concept of *evolution*.

The first proposals in this direction date back to the mid sixties, when in the United States John Holland, at the University of Michigan, introduced Genetic Algorithms (Holland 1975, Goldberg 1989, Davis 1991) and Lawrence Fogel, of the University of California at San Diego, came up with Evolutionary Programming (Fogel et al. 1966), while in Europe Ingo Rechenberg and Hans-Paul Schwefel, at the Technical University of Berlin, independently began to work on Evolution Strategies (Rechenberg 1973, Schwefel 1981). Their pioneering works eventually gave rise to a broad class of optimization methods particularly well suited for hard problems where little is known about the underlying search space.
The last development of this research thread is the so called Genetic Programming, introduced by John Koza (1993), at the Stanford University.
A recent work of reference and synthesis in the field of Evolutionary Algorithms is Michalewicz (1992).
An evolutionary algorithm maintains a population of candidate solutions for the problem at hand, and makes it evolve by iteratively applying a (usually quite small) set of stochastic operators, known as *mutation*, *recombination*, *reproduction* and *selection*.
Mutation randomly perturbs a candidate solution; recombination decomposes two distinct solutions and then randomly mixes their parts to form a novel solution; reproduction replicates the most successful solutions found in a population at a rate proportional to their relative goodness; while selection purges poor solutions from a population.
The initial population may be either a random sample of the solution space or may be seeded with solutions found by simple local search procedures, if these are available.
The resulting process tends to find globally optimal solutions to the problem much in the same way as in Nature populations of organisms tend to adapt to their surrounding environment.

4. Problem formulation

Portfolio optimization can be viewed as a convex combination parametric programming problem (Steuer 1986), in that an investor wants to minimize risk while maximizing expected return; on the efficient frontier a larger expected return corresponds to a greater risk. This can be expressed as a two-objective optimization problem:

$$\min_{\mathbf{w}}\{\text{Risk}(\mathbf{w})\}$$
$$\max_{\mathbf{w}}\{\text{Return}(\mathbf{w})\}$$

subject to

$$\sum_i w_i = 1;$$
$$w_i \geq 0.$$

These two objectives can be parametrized to yield a convex combination parametric programming problem with objective

$$\min_{\mathbf{w}}\{\lambda\,\text{Risk}(\mathbf{w}) - (1-\lambda)\,\text{Return}(\mathbf{w})\}$$

where parameter λ is a trade-off coefficient ranging between 0 and 1. When $\lambda = 0$ the investor disregards risk and only seeks to maximize expected return; when $\lambda = 1$ the risk
alone is being minimized, whatever the expected return.
Since there is no general way to tell which particular trade-off between risk and return is to be considered the best, optimizing a portfolio means finding a whole range of optimal portfolios for all the possible values of the trade-off coefficient; the investors will thus be able to choose the one they believe appropriate for their requirements.

5. An evolutionary approach

The structure of the problem itself suggests a global architecture for an evolutionary algorithm, many solutions being actually needed—one for each trade-off coefficient considered. A natural way to achieve that in an evolutionary setting is to have a number of distinct populations evolve in parallel with different trade-off coefficient values. The greater this number, the finer the resolution with which an investor will be able to examine the efficient frontier.
Because it is likely that slight variations in the trade-off coefficient do not significantly worsen a good solution, a natural way to sustain the evolutionary process is to allow migration or cross-breeding between individuals belonging to immediately neighboring populations.

5.1. Algorithm Overview

This idea has been implemented by breaking up the population into several subpopulations, or *species,* one for each value of the trade-off coefficient; individuals from a species are allowed to mate only with individuals from the same or an adjacent species, and new individuals always replace the worst individual of their own species, so that the balance among species is always maintained.
The evolutionary algorithm is based on a modified version of Darrel Whitley's GENITOR system (Whitley and Kauth 1988), which uses a steady-state breeding strategy and elitist selection.
The overall structure of the algorithm is as follows:

create initial population;
repeat
    select(γ);
    select(κ) **such that** $\left|\lambda_\kappa - \lambda_\gamma\right| \leq 1$;
    ζ := crossover(γ, κ);
    mutate(ζ);
    replace(worst(λ_ζ), ζ);
until termination condition.

5.2. Encoding and fitness

As pointed out by Michalewicz and Janikow (1992), the most effective strategy for handling constraints in evolutionary algorithms consists, as long as this is feasible, in choosing an appropriate encoding and reproduction procedure such that illegal solutions

are never generated or, even better, such that any possible chromosome encodes for a legal solution. This relieves the algorithm from having to check new individuals for constraint satisfaction.
The encoding that has been selected is intuitive and has the advantage that only legal solutions can be represented: the share of asset *i* is encoded in the genotype γ as an integer γ_i between 0 and a positive constant $g_{\max}$. The actual weight w_i of that share is computed as the ratio of its encoded value to the total sum of all the encoded values in the chromosome,

$$w_i = \gamma_i \Big/ \sum_j \gamma_j .$$

An additional integer γ_0 encodes the trade-off coefficient λ to be used to evaluate the objective function $z(\mathbf{w}; \lambda)$ for the corresponding portfolio **w**.
The objective function is defined as

$$z(\mathbf{w};\lambda) = (1-\lambda)\,\mathrm{Return}(\mathbf{w}) - \lambda\,\mathrm{DSR}(\mathbf{w}),$$

where the return and downside-risk are appropriately scaled.
In GENITOR a lower fitness value corresponds to a better solution; DRAGO scales fitness in such a way that the best individual of each species always has zero fitness. The fitness of individual γ, corresponding to portfolio **w**, is calculated as

$$f(\gamma) = \frac{z_{\max}(\lambda_\gamma) - z(\mathbf{w};\lambda_\gamma)}{z_{\max}(\lambda_\gamma) - z_{\min}(\lambda_\gamma) + 1},$$

where $z_{\max}(\lambda)$ and $z_{\min}(\lambda)$ stand for the maximum and minimum values of the objective function for the best and worst individuals respectively in the species associated to the trade-off coefficient λ.

5.3. Genetic operators

A variation of uniform crossover has been employed as a genetic recombination operator suitable for the problem: let γ and κ be two parent chromosomes and

$$\mathrm{TOT}_\gamma = \sum_i \gamma_i , \qquad \mathrm{TOT}_\kappa = \sum_i \kappa_i ;$$

furthermore, suppose that the *i*th gene in γ, γ_i has been chosen to be passed onto the offspring: then the value of the *i*th gene in the offspring ζ is

$$\zeta_i = \min\left(g_{\max}, \gamma_i \frac{\mathrm{TOT}_\gamma + \mathrm{TOT}_\kappa}{2\,\mathrm{TOT}_\gamma} \right).$$

The mutation operator adopted by DRAGO alters a gene either by incrementing or decrementing it by one. This ensures that a chromosome undergoing mutation will not experience abrupt changes (mutation should not disrupt schemata).
While preserving schemata, this mutation strategy still involves a certain randomness, for it implies that a chromosome with a high total sum is less sensitive to mutations,

whereas a chromosome with a small total sum feels the effect of a mutation more strongly.

6. Experiments and Results

Since its implementation, DRAGO has been extensively exploited by SIGE Consulenza to optimize some core portfolios of the mutual funds offered to the public and it has thus undergone a comprehensive testing on the field, with portfolios ranging from a few dozens to hundreds of assets.
To demonstrate the typical working of the algorithm, Fig. 3 shows snapshots of the evolution of a population of 200 individuals over 10000 "generations" (here generation means the birth of a new individual and the death of the worst one). The population is segmented into 20 species, the mutation rate is set to 0.5 and the risk aversion coefficient q is set to 2, i.e. target semivariance is the index of risk.
Note that the space of risk and return is not the original problem space (which in this case has 15 dimensions), but only a convenient transformed space to view solutions; therefore, although the initial population is uniformly distributed over the original problem space, it does not look so in the figure.

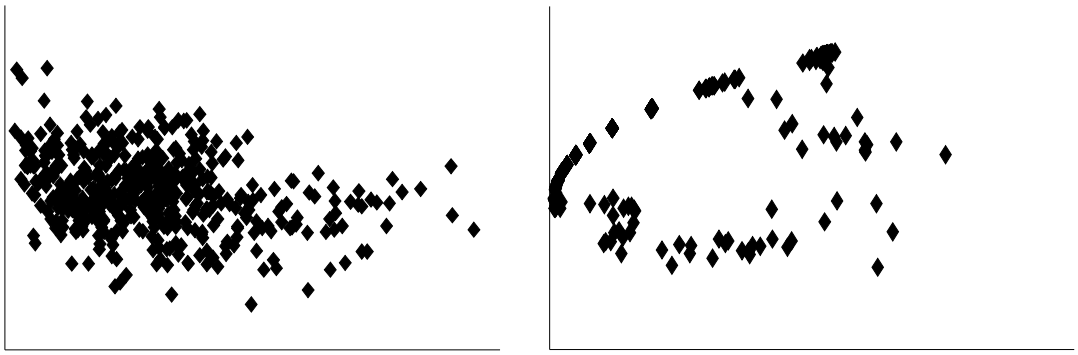

Fig. 3

Execution times to obtain satisfactory solutions range between 5 minutes for the simplest problems and 10 hours for more complex cases on a quite ordinary 25MHz 486DX-based personal computer.

7. Conclusion

The use of an evolutionary approach to portfolio optimization has made possible to employ in an operational setting meaningful indices of risk other than variance, which would not be available with conventional methods.
The resulting decision support system shows response times not amazing on an absolute scale, but adequate for most practical uses, especially because the average portfolio manager deals with typical investment horizons in the order of months, so that precise timing is not an issue.

References

Bawa V.S.: Optimal rules for ordering uncertain prospects. *Journal of Financial Economics*, March 1975.

Davis L.: *Handbook of genetic algorithms*. VNR Computer Library, Van Nostrand Reinhold, New York, 1991.

Fogel L.J., Owens A.J. and Walsh M.J.: *Artificial intelligence through simulated evolution*. John Wiley & Sons, New York, 1966.

Goldberg D.E.: *Genetic algorithms in search, optimization & machine learning*. Addison Wesley, Reading, MA, 1989.

Harlow H.V.: Asset allocation in a downside-risk framework. *Financial Analysts Journal*, September/October 1991, 30-40.

Holland J.H.: *Adaptation in Natural and Artificial Systems*, MIT Press, Cambridge, MA, 1975.

Koza J.R.: *Genetic Programming: on the programming of computers by means of natural selection*, MIT Press, Cambridge, MA, 1993.

Markowitz H.M.: *Portfolio selection*. John Wiley & Sons, New York, 1959.

Micalewicz Z.: *Genetic algorithms + data structures = evolution programs*. Spinger-Verlag, Berlin, 1992.

Michalewicz Z. and Janikow C.Z.: Handling constraints in genetic algorithms. In: *Proceedings of the Fourth International Conference on Genetic Algorithms*, 1992.

Rechenberg I.: *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Fromman-Holzboog Verlag, Stuttgart, 1973.

Schwefel H.-P.: *Numerical optimization of computer models,* John Wiley & Sons, New York, 1981.

Steuer R.E.: *Multiple criteria optimization: Theory, computation and application*. John Wiley & Sons, New York, 1986.

Whitley D. and Kauth J.: *GENITOR: a different genetic algorithm*. Technical Report CS-88-101, Colorado State University, Fort Collins, CO, 1988.