Swarm and Evolutionary Computation ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

Swarm and Evolutionary Computation

journal homepage: www.elsevier.com/locate/swevo

Regular Paper

Evolution on trees: On the design of an evolution strategy for scenario-based multi-period portfolio optimization under transaction costs

Hans-Georg Beyer, Steffen Finck*, Thomas Breuer

FH Vorarlberg University of Applied Sciences, Hochschulstrasse 1, 6850 Dornbirn, Austria

A R T I C L E I N F O

Article history:
Received 23 May 2013
Received in revised form
11 February 2014
Accepted 19 March 2014

Keywords:
Evolution strategy
Covariance matrix adaptation
Portfolio optimization
Self-adaptation

A B S T R A C T

Scenario-based optimization is a problem class often occurring in finance, planning and control. While the standard approach is usually based on linear stochastic programming, this paper develops an Evolution Strategy (ES) that can be used to treat nonlinear planning problems arising from Value at Risk (VaR)-constraints and not necessarily proportional transaction costs. Due to the VaR-constraints the optimization problem is generally of non-convex type and its decision version is already NP-complete. The developed ES is the first algorithm in the field of evolutionary and swarm intelligence that tackles this kind of optimization problem. The algorithm design is based on the covariance matrix self-adaptation ES (CMSA-ES). The optimization is performed on scenario trees where in each node specific constraints (balance equations) must be fulfilled. In order to evaluate the performance of the ES proposed, instances of increasing problem hardness are considered. The application to the general case with nonlinear node constraints shows not only the potential of the ES designed, but also its limitations. The latter are basically determined by the high dimensionalities of the search spaces defined by the scenario trees.

1. Introduction

This paper designs evolution strategies (ESs) for discrete-time multi-period multi-asset portfolio optimization problems with Value at Risk (VaR)-constraints and not necessarily proportional transaction costs. It is a follow-up of [1], which analyzed the design of Evolution Strategies (ES) for one-period portfolio optimization problems under VaR-constraints without transaction costs, where it turned out that the algorithm design was challenging due to the combination of seemingly simple constraints. The problem considered in this paper is even more complex since in the multi-period case, node constraints (in general non-linear balance equations) must be fulfilled forcing the ES to evolve on a non-linear manifold. The considered problem is a stochastic control problem whose information and decision structure is defined on scenario trees. In finance and operations research literature, problems of this kind are approached by linear stochastic programming [2,3]. However, due to the nonlinearities in the problem class considered here, the problem must be linearized to allow for the application of this standard approach. In [4] it is stated that metaheuristics might be successfully applied to such types of problems. However, up to now this problem class has not been tackled by swarm or evolutionary methods. Since Evolutionary Algorithms (EAs) allow for a great flexibility they might be well-suited for the treatment of nonlinearities arising from VaR-constraints and non-linear transaction costs.

In contrast to the problem considered in this paper, classical portfolio optimization in the framework of the Markowitz model is intrinsically a multiobjective optimization problem. In [5] recent trends for EAs applied to such portfolio optimization problems are discussed and in [6] a particle swarm approach has been proposed and compared to other algorithms to tackle that problem class.

However, virtually all the featured algorithms found in the literature deal only with the one-period problem and the multi-period problem is not considered. In cases where an EA was applied to the multi-period portfolio problem [7–9], the problem itself is different from (and to some extent simpler) the one considered here.

This paper is the first that provides a *design methodology* and an ES implementation for evolutionary optimization on scenario trees encountered in non-linear multi-period stochastic problems. Section 2 introduces the problem in its general form. Section 3 deals with the problem in the special case of vanishing transaction

* Corresponding author. Tel.: +43 5572 7927122; fax: +43 5572 7929510.
E-mail addresses: Hans-Georg.Beyer@fhv.at (H.-G. Beyer), steffen.finck@fhv.at (S. Finck), Thomas.Breuer@fhv.at (T. Breuer).

http://dx.doi.org/10.1016/j.swevo.2014.03.002
2210-6502/

costs and uses the standard CMSA-ES as an algorithmic skeleton to design an ES, the multi-period (mp) CMSA-ES, that operates on the tree structure of the optimization problem. Section 4 presents an experimental evaluation of the novel mpCMSA-ES. While the development of Section 3 is based on linear balance equations in the tree nodes, Section 5 extends the approach to the problem in its general form with nonlinear constraints arising for example from non-linear transaction costs. Section 6 summarizes the paper and provides an outlook.

2. The optimization problem

In many financial optimization models the uncertainty in asset prices (or other risk factors) is represented by a number of mass points forming a scenario tree. Such trees are often interpreted as approximating stochastic processes, but interpretational issues involved in such approximations will be of no concern here. Consider an ordered directed tree, the nodes of which are indexed by $n \in \{0, \dots, N\}$. The zero node is the root node. The predecessor node k of a node n is indexed by $k = \pi(n)$ (representing an adjacency list). There are N_L leaf nodes representing the terminal nodes having no child nodes. For sake of simplicity the leaf nodes are indexed by the consecutive numbers in $\mathbb{N}_L = \{N - N_L + 1, \dots, N\}$. For each leaf node there is a unique path leading to it. Denote by p_l, $l \in \mathbb{N}_L$, the probability of arriving at leaf node l.[1]

With various methods of scenario reduction the size of the scenario tree can be reduced [10,11]. This is essential for numerical tractability, but in case the tree is too sparse arbitrage cannot be ruled out [12,13]. A short selling constraint like (2) below still ensures the existence of a solution to the portfolio optimization problem, but this solution will in general be biased.

Before transaction costs, the wealth W in a node n is[2]

$$W_n := \boldsymbol{x}_n^{\mathrm{T}} \boldsymbol{\xi}_n = \sum_{m=1}^{M} (\boldsymbol{x}_n)_m (\boldsymbol{\xi}_n)_m, \tag{1}$$

where $\boldsymbol{x} = (x_1, \dots, x_M)^{\mathrm{T}}$ is the M-dimensional vector of the nodal portfolio weights x_m

$$\forall m = 1, \dots, M : x_m \ge 0, \tag{2}$$

which amounts to a short selling constraint. The vector

$$\boldsymbol{\xi} = (\xi_1, \dots, \xi_M)^{\mathrm{T}}, \quad \forall m = 1, \dots, M: \ \xi_m > 0 \tag{3}$$

contains the M asset prices in the respective node. Starting with an initial capital W_0 in the root node and the asset prices $\boldsymbol{\xi}_0$

$$W_0 = \boldsymbol{x}_0^{\mathrm{T}} \boldsymbol{\xi}_0, \tag{4}$$

the decision maker has to choose the portfolio weights $\boldsymbol{x}_0$ at time zero (being in the root node). At the next time step, a number of child scenarios with changed asset prices $\boldsymbol{\xi}_k$ are considered. As a result of asset price changes in the transition from node $\pi(k)$ to node k, the value of the portfolio will change to $W_k = \boldsymbol{x}_{\pi(k)}^{\mathrm{T}} \boldsymbol{\xi}_k$ in the child node k. Now the decision maker can change the nodal portfolio weights in the child node constrained by the available wealth W_k

$$\forall k = 1, \dots, N - N_L : \boldsymbol{x}_k^{\mathrm{T}} \boldsymbol{\xi}_k = \boldsymbol{x}_{\pi(k)}^{\mathrm{T}} \boldsymbol{\xi}_k. \tag{5}$$

Additionally, one can take transaction costs [14–19] into account. The respective cost function will be "B[uy]S[ell]". A simple option is

$$\mathrm{BS}_1(\boldsymbol{x}_\pi, \boldsymbol{x}, \boldsymbol{\xi}) := c \sum_{m=1}^{M} |(\boldsymbol{x}_\pi)_m - (\boldsymbol{x})_m| (\boldsymbol{\xi})_m \tag{6}$$

which accounts for buy and sell transaction costs in a symmetric and proportional way, assuming vanishing fixed costs of transactions. A somewhat more elaborated model might consider fixed costs $c_f \ge 0$ and different costs for buying $c_b \ge 0$ and selling $c_s \ge 0$

$$\begin{aligned} \mathrm{BS}_2(\boldsymbol{x}_\pi, \boldsymbol{x}, \boldsymbol{\xi}) := {} & c_f \sum_{m=1}^{M} \overline{\delta}((\boldsymbol{x}_\pi)_m - (\boldsymbol{x})_m) \\ & + c_b \sum_{m=1}^{M} \Theta((\boldsymbol{x})_m - (\boldsymbol{x}_\pi)_m) |(\boldsymbol{x}_\pi)_m - (\boldsymbol{x})_m| (\boldsymbol{\xi})_m \\ & + c_s \sum_{m=1}^{M} \Theta((\boldsymbol{x}_\pi)_m - (\boldsymbol{x})_m) |(\boldsymbol{x}_\pi)_m - (\boldsymbol{x})_m| \ (\boldsymbol{\xi})_m, \end{aligned} \tag{7}$$

where $\overline{\delta}$ is the anti-delta function and Θ is the step function, as defined in (A.9) and (A.10) in the Appendix.

Now, the decision maker may change the weights in the child nodes such that the wealth after a transaction equals the wealth before transaction minus the transaction costs. That is, in each inner node (i.e. those nodes not being root node or terminal nodes) a equality constraint must hold:

$$\forall k = 1, \dots, N - N_L : \boldsymbol{x}_k^{\mathrm{T}} \boldsymbol{\xi}_k + \mathrm{BS}(\boldsymbol{x}_{\pi(k)}, \boldsymbol{x}_k, \boldsymbol{\xi}_k) = \boldsymbol{x}_{\pi(k)}^{\mathrm{T}} \boldsymbol{\xi}_k. \tag{8}$$

The last portfolio choices are made at the predecessors of the leaf nodes. At the leaf nodes themselves a final wealth value is realized.

Thus a trading strategy $\mathbf{X}$ is described by a matrix made of the $N - N_L + 1$ non-terminal node decision vectors $\boldsymbol{x}_k$ (being column vectors)

$$\mathbf{X} = (\boldsymbol{x}_0, \dots, \boldsymbol{x}_{N - N_L}), \tag{9}$$

which satisfy the budget constraints (8). The expected final wealth is $\sum_{l = N - N_L + 1}^{N} p_l W_l$, where $W_l = \boldsymbol{x}_{\pi(l)}^{\mathrm{T}} \boldsymbol{\xi}_l$. The objective is to maximize this expected wealth at the end of the planning horizon:

$$f(\mathbf{X}) := \sum_{l = N - N_L + 1}^{N} p_l \ \boldsymbol{x}_{\pi(l)}^{\mathrm{T}} \boldsymbol{\xi}_l \longrightarrow \max_{\boldsymbol{x}_0, \dots, \boldsymbol{x}_{N - N_L}}. \tag{10}$$

In addition to the budget constraints (8) a risk constraint on the admissible trading strategies will be imposed. Regulators [20] require banks to hold at the end of the trading day sufficient economic capital defined in terms of Value at Risk (VaR).[3] These regulatory VaR-constraints translate into the requirement that at least an α-fraction of final wealth values W_l is greater than (or equal to) a κ-multiple of the initial capital W_0. The function $\eta(\mathbf{X})$ sums all p_l for which this holds

$$\eta(\mathbf{X}) := \left\{ \sum_l p_l | l \in \mathbb{N}_L \wedge W_l \ge \kappa W_0 \right\}, \tag{11}$$

resulting in the nonlinear inequality constraint:

$$\eta(\mathbf{X}) \ge \alpha. \tag{12}$$

Let us summarize the optimization problem

$$\sum_{l = N - N_L + 1}^{N} p_l \boldsymbol{x}_{\pi(l)}^{\mathrm{T}} \boldsymbol{\xi}_l \longrightarrow \max_{\mathbf{X}}, \tag{13a}$$

s. t.

[1] Equivalently one could specify conditional probabilities for each edge of the tree. The conditional probabilities of each edge leaving the same node add up to one. In this formulation, p_l equals the product of conditional probabilities of all edges along the path leading from the root node to the leaf node l.

[2] Note, the notation $(\cdot)_m$ refers to the mth component of a vector, i.e., $(\boldsymbol{x})_m \equiv x_m$.

[3] Determining capital requirements defined in terms of some coherent or at least convex risk measure would be more satisfactory, see [21,22]. Actually, a constraint in the coherent risk measure CVaR would be easier to handle than a VaR-constraint, since it leads to a linear problem [23]. Furthermore, it would be more satisfactory to have a full-fledged capital requirement for processes rather than for the final wealth at the end of the trading period, since new information and transactions arise during the trading period [3,24–26]. The current regulatory framework is taken as given.