Global Optimization of the Scenario Generation and Portfolio Selection Problems

Panos Parpas and Berç Rustem

Department of Computing, Imperial College, London SW7 2AZ

Abstract. We consider the global optimization of two problems arising from financial applications. The first problem originates from the portfolio selection problem when high-order moments are taken into account. The second issue we address is the problem of scenario generation. Both problems are non-convex, large-scale, and highly relevant in financial engineering. For the two problems we consider, we apply a new stochastic global optimization algorithm that has been developed specifically for this class of problems. The algorithm is an extension to the constrained case of the so called diffusion algorithm. We discuss how a financial planning model (of realistic size) can be solved to global optimality using a stochastic algorithm. Initial numerical results are given that show the feasibility of the proposed approach.

1 Introduction

We consider the global optimization of two problems arising from financial applications. The first problem originates from the portfolio selection problem when high-order moments are taken into account. This model is an extension of the celebrated mean-variance model of Markowitz[1, 2]. The inclusion of higher order moments has been proposed as one possible augmentation to the model in order to make it more applicable. The applicability of the model can be broadened by relaxing one of its major assumptions, i.e. that the rate of returns are normal. The second issue we address is the problem of scenario generation i.e. the description of the uncertainties used in the portfolio selection problem. Both problems are non-convex, large-scale, and highly relevant in financial engineering.

Given the numerical and theoretical challenges presented by these models we only consider the "vanilla" versions of the two problems. In particular, we focus on a single period model where the decision maker (DM) provides as input preferences with respect to mean, variance, skewness, and possibly kurtosis of the portfolio. Using these four parameters we then formulate the multi-criteria optimization problem as a standard nonlinear programming problem. This version of the decision model is a non-convex linearly constrained problem.

Before we can solve the portfolio selection problem we need to describe the uncertainties regarding the returns of the risky assets. In particular we need to specify: (1) the possible states of the world and (2) the probability of each state. A common approach to this modeling problem is the method of matching moments (see e.g. [3, 4, 5]). The first step in this approach is to use the historical

M. Gavrilova et al. (Eds.): ICCSA 2006, LNCS 3982, pp. 908–917, 2006.

data in order to estimate the moments (in this paper we consider the first four central moments i.e. mean, variance, skewness, and kurtosis). The second step is to compute a discrete distribution with the same statistical properties as the ones calculated in the previous step. Given that our interest is on real–world applications we recognize that there may not always be a distribution that matches the calculated statistical properties. For this reason we formulate the problem as a least squares problem [3, 4]. The rationale behind this formulation is that we try to calculate a description of the uncertainty that matches our beliefs as well as possible. The scenario generation problem also has a non-convex objective function, and is linearly constrained.

For the two problems described above we apply a new stochastic global optimization algorithm that has been developed specifically for this class of problems. The algorithm is described in [6] (see also section 4). It is an extension to the constrained case of the so called diffusion algorithm [7, 8, 9, 10]. The method follows the trajectory of an appropriately defined Stochastic Differential Equation (SDE). Feasibility of the trajectory is achieved by projecting its dynamics onto the set defined by the linear equality constraints. A barrier term is used for the purpose of forcing the trajectory to stay within any bound constraints (e.g. positivity of the probabilities, or bounds on how much of each asset to own).

The purpose of this paper is to show that stochastic optimization algorithms can be used to solve realistic financial planning problems. A review of applications of global optimization to portfolio selection problems appeared in [11]. A deterministic global optimization algorithm for a multi-period model appeared in [12]. This paper extends and complements the methods mentioned above in the sense that we describe a complete framework for the solution of a realistic financial model. The type of models we consider, due to the large number of variables, cannot be solved by deterministic algorithms. Consequently, practitioners are left with two options: solve a simpler, but less relevant model, or use a heuristic algorithm (e.g. tabu–search or evolutionary algorithms). The approach proposed in this paper lies somewhere in the middle. The proposed algorithm belongs to the simulated–annealing family of algorithms, and it has been shown in [6] that it converges to the global optimum (in a probabilistic sense). Moreover, the computational experience reported in [6] seems to indicate that the method is robust (in terms of finding the global optimum) and reliable. We believe that such an approach will be useful in many practical applications. Admittedly the models (especially the portfolio selection problem) are rather simplistic. Given the theoretical and computational difficulties involved with such models it is important to consider the simplified version of the problem in the hope that this approach will shed more light to the general case. Moreover, to the authors' knowledge this is the first paper to address, in a holistic manner, the global optimization of the moment problem and the portfolio selection problem with higher order moments.

The rest of the paper is structured as follows: in section 2 we describe the scenario generation problem. While there are many ways to generate scenario trees for stochastic programming problems, we will focus on the moment matching approach. The interested reader is referred to [3] for a review of other methods. Also in this section we discuss the importance of arbitrage opportunities; we describe how we dealt with this requirement of financial models in our implementation. In section 3 we discuss the portfolio selection problem. A model with a non-convex objective and linear constraints is proposed as a simple extension to the classical Markowitz model. The non-convexities in the model arise from the inclusion of higher order moments. The model considered here relaxes the normality assumption of the classical model, the reader is referred to [13] for a more complete overview of non-convex optimization problems in financial applications. In section 4 we describe an algorithm for the solution of the two models described above. For a full theoretical treatment of the algorithm we refer the interested reader to [6]. In section 5 we present some initial numerical experiments. We study how difficult (in terms of computation time) it is to compute an arbitrage free scenario tree. We also study how the global optimum changes as we vary the parameters of the model. To illustrate the effect of the parameters we present some 3-dimensional plots of efficient frontiers, the analogs of the classical Markowitz efficient frontiers.

2 Scenario Generation

From its inception Stochastic Programming (SP) has found several diverse applications as an effective paradigm for modeling decisions under uncertainty. The focus of initial research was on developing effective algorithms for models of realistic size. An area that has only recently received attention is on methods to represent the uncertainties of the decision problem.

A review of available methods to generate meaningful descriptions of the uncertainties from data can be found in[3]. We will use a least squares formulation (see e.g. [3, 4]). It is motivated by the practical concern that the moments, given as input, may be inconsistent. Consequently the best one can do is to find a distribution that fits the available data as well as possible. It is further assumed that the distribution is discrete. Under these assumptions the problem can be written as:

$$\min_{\omega,p} \sum_{i=1}^{n} \Big(\sum_{j=1}^{k} p_j m_i(\omega_j) - \mu_i \Big)^2$$
$$\text{s.t} \quad \sum_{j=1}^{k} p_j = 1 \quad p_j \geq 0 \quad j = 1, \ldots, k$$

Where μ_i represent the statistical properties of interest, and $m_i(\cdot)$ is the associated 'moment' function. For example, if μ_i is the target mean for the i^{th} asset then $m_i(\omega_j) = \omega_j^i$ i.e. the j^{th} realization of the i^{th} asset. Numerical experiments using this approach for a multistage model, were reported in [4] (without arbitrage considerations). Other methods such as maximum entropy[14], and semidefinite programming [15], while they enjoy strong theoretical properties they cannot be used when the data of the problem are inconsistent. A disadvantage of the least squares model is that it is highly non-convex which makes it very difficult to handle numerically. These considerations have lead to the development of the algorithm described in section 4 (see also [6]) that can efficiently compute global optima for problems in this class.

When using scenario trees for financial planning problems it becomes necessary to address the issue of arbitrage opportunities[4, 16]. An arbitrage opportunity is a self–financing trading strategy that generates a strictly positive cash flow in at least one state and whose payoffs are nonnegative in all other states. In other words it is possible to get something for nothing. In our implementation we eliminate arbitrage opportunities by computing a sufficient set of states so that the resulting scenario tree has the arbitrage free property. This is achieved by a simple two step process. In the first step we generate random rates of returns, these are sampled by a uniform distribution. We then test for arbitrage by solving the system:

$$x_0^i = e^{-r} \sum_{j=1}^{m} x_j^i \pi_j \ , \sum_{j=1}^{m} \pi_j = 1, \ , \pi_j \geq 0, \quad j = 1, \ldots, m \ i = 1, \ldots, n. \tag{1}$$

Where x_0^i represents the current (known) state of the world for the i^{th} asset, x_j^i represents the j^{th} realization of the i^{th} asset in the next time period (these are generated by the simulations mentioned above). r is the risk-less rate of return. The π_j are called the risk neutral probabilities. According to a fundamental result of Harisson and Kerps [17], the existence of the risk neutral probabilities is enough to guarantee that the scenario tree has the desired property. In the second step, we solve the least squares problem with some of the states fixed to the states calculated in the first step. In other words, we solve the following problem:

$$\begin{aligned} \min_{\omega,p} & \sum_{i=1}^{n} \Big(\sum_{j=1}^{k} p_j m_i(\omega_j) + \sum_{l=1}^{m} p_l m_i(\hat{\omega}_l) - \mu_i \Big)^2 \\ \text{s.t} & \sum_{j=1}^{k+m} p_j = 1 \quad p_j \geq 0 \quad j = 1, \ldots, k+m \end{aligned} \tag{2}$$

In the problem above, $\hat{\omega}$ are fixed. Solving the preceding problem guarantees a scenario tree that is arbitrage free.

3 Portfolio Selection

In this section we describe the portfolio selection problem when higher order terms are taken into account. The classical mean–variance approach to portfolio analysis seeks to balance risk (measured by variance) and reward (measured by expected value). There are many ways to specify the single period problem. We will be using the following basic model:

$$\min_{w} -\alpha\mathbb{E}[w] + \beta\mathbb{V}[w]$$
$$\text{s.t} \sum_{i=1}^{n} w_i = 1 \quad l_i \leq w_i \leq u_i \; i = 1, \ldots, n. \tag{3}$$

Where $\mathbb{E}[\cdot]$ and $\mathbb{V}[\cdot]$ represent the mean rate of return and its variance respectively. The single constraint is known as the *budget constraint* and it specifies the initial wealth (without loss of generality we have assumed that this is one). The α and β are positive scalars, and are chosen so that $\alpha + \beta = 1$. They specify the DMs preferences, i.e. $\alpha = 1$ means that the DM is risk–seeking, while $\beta = 1$ implies that the DM is risk averse. Any other selection of the parameters will produce a point on the efficient frontier. The decision variable (w) represents the commitment of the DM to a particular asset. Note that this problem is a convex quadratic programming problem for which very efficient algorithms exists. The interested reader is referred to the review in [13] for more information regarding the Markowitz model.

We propose an extension of the mean–variance model using higher order moments. The vector optimization problem can be formulated as a standard non-convex optimization problem using two additional scalars to act as weights. These weights are used to enforce the DMs preferences. The problem is then formulated as follows:

$$\min_{w} -\alpha\mathbb{E}[w] + \beta\mathbb{V}[w] - \gamma\mathbb{S}[w] + \delta\mathbb{K}[w]$$
$$\text{s.t} \sum_{i=1}^{n} w_i = 1 \quad l_i \leq w_i \leq u_i \; i = 1, \ldots, n. \tag{4}$$

Where $\mathbb{S}[\cdot]$ and $\mathbb{K}[\cdot]$ represent the skewness and kurtosis of the rate of return respectively. γ and δ are positive scalars. The four scalar parameters are chosen so that they sum to one. Positive skewness is desirable (since it corresponds to higher returns albeit with low probability) while kurtosis is undesirable since it implies that the DM is exposed to more risk. The model in (4) can be extended to multiple periods while maintaining the same structure (non convex objective and linear constraints). The numerical solution of (2) and (4) will be discussed in the next two sections.

4 A Stochastic Optimization Algorithm

The models described in the previous section can be written as:

$$\min_{x} f(x)$$
$$\text{s.t } Ax = b$$
$$x \geq 0.$$

A well known method for obtaining a solution to an unconstrained optimization problem is to consider the following Ordinary Differential Equation (ODE):

$$dX(t) = -\nabla f(X(t))dt. \tag{5}$$

By studying the behavior of $X(t)$ for large t, it can be shown that $X(t)$ will eventually converge to a stationary point of the unconstrained problem. A review of, so called, continuous-path methods can be found in [18]. A deficiency of using (5) to solve optimization problems, is that it will get trapped in local minima. In order to allow the trajectory to escape from local minima, it has been proposed by various authors (e.g. [7, 8, 9, 10]) to add a stochastic term that would allow the trajectory to "climb" hills. One possible augmentation to (5) that would enable us to escape from local minima is to add noise. One then considers the *diffusion process*:

$$dX(t) = -\nabla f(X(t))dt + \sqrt{2T(t)}dB(t). \tag{6}$$

Where $B(t)$ is the standard Brownian motion in $\mathbb{R}^n$. It has been shown in [8, 9, 10], under appropriate conditions on f, and $T(t)$, that as $t \to \infty$, the transition probability of $X(t)$ converges (weakly) to a probability measure Π. The latter, has its support on the set of global minimizers.

For the sake of argument, suppose we did not have any linear constraints, but only positivity constraints. We could then consider enforcing the feasibility of the iterates by using a barrier function. According to the algorithmic framework sketched-out above, we could obtain a solution to our (simplified) problem, by following the trajectory of the following SDE:

$$dX(t) = -\nabla f(X(t))dt + \mu X(t)^{-1}dt + \sqrt{2T(t)}dB(t). \tag{7}$$

Where $\mu > 0$, is the barrier parameter. By X^{-1}, we will denote an n-dimensional vector whose i^{th} component is given by $1/X_i$. Having used a barrier function to deal with the positivity constraints, we can now introduce the linear constraints into our SDE. This process has been carried out in [6] using the projected SDE:

$$dX(t) = P[-\nabla f(X(t)) + \mu X(t)^{-1}]dt + \sqrt{2T(t)}PdB(t). \tag{8}$$

Where, $P = I - A^T(AA^T)^{-1}A$. The proposed algorithm works in a similar manner to gradient projection algorithms. The key difference is the addition of a barrier parameter for the positivity of the iterates, and a stochastic term that helps the algorithm escape from local minima.

The global optimization problem can be solved by fixing μ, and following the trajectory of (8) for a suitably defined function $T(t)$. After sufficiently enough time passes, we reduce μ, and repeat the process. The proof that following the trajectory of (8) will eventually lead us to the global minimum appears in [6]. Note that the projection matrix for the type of constraints we need to impose for our models is particularly simple. For a constraint of the type: $\sum_{i=1}^n x_i = 1$ the projection matrix is given by:

$$P_{ij} = \begin{cases} -\frac{1}{n} & \text{if } i \neq j, \\ \frac{n-1}{n} & \text{otherwise.} \end{cases}$$

5 Numerical Experiments

The algorithm described in the previous section was implemented in C++. Before we discuss our numerical results we provide some useful implementation details. From similar studies in the unconstrained case[7] and box constrained case[19], we know that a deficiency of stochastic methods (of the type proposed in this paper) is that they require a large number of function evaluations. The reason of this shortcoming is that the annealing schedule has to be sufficiently slow in order to allow the trajectory to escape from local minima. Therefore, whilst there are many sophisticated methods for the numerical solution of SDEs, we decided to use the cheaper stochastic Euler method. The latter method is a generalization of the well known Euler method, for ODEs, to the stochastic case. The main iteration is given by:

$$X(t+1) = X(t) + P[-\nabla f(X(t)) + \mu X(t)^{-1}]\Delta t + \sqrt{2T(t)\Delta t}Pu.$$

Where Δt is the discritized step length parameter, u is a standard Gaussian vector, i.e. $u \sim N(0, I)$, and $X(0)$ is chosen to be strictly feasible.

The algorithm starts by dividing the discritized time into k periods. Following a single trajectory will be too inefficient. Therefore, starting from a single strictly feasible point the algorithm generates m different trajectories. After a single period elapses, we remove the worst performing trajectory. Since, all trajectories generate feasible points, we can assess the quality of the trajectory by the best objective function value achieved on the trajectory. We then randomly select one of the remaining trajectories, and duplicate it. At this stage we reduce the noise coefficient of the duplicated trajectory.

When all the periods have been completed, in the manner described above, we count this event as one iteration. At this point we reduce the barrier parameter. This parameter is started at $\mu = 0.1$, and reduced by 0.75 at every iteration. We then repeat the same process, with all the trajectories starting from the best point found so far. If the current incumbent solution vector remained the same for more than l iterations ($l > 4$, in our implementation) then we reset the noise to its initial value. The algorithm terminates when the noise term is smaller than a predefined value ($0.1e-4$) or when after five successive resets of the noise term, no improvement could be made. In our implementation we used two trajectories, two periods (each of length $20e4$). We used an initial value of 10 for the annealing schedule, and reduced it by 0.6 at every iteration. The same parameters were used for all the simulations.

In table 1 we show the computational effort required to compute the sufficient set of states required to guarantee the arbitrage free property of the scenario tree. The numbers shown are averages of 50 runs. It is clear from table 1 that it is relatively easy to find the sufficient states. However, as the number of assets increases the number of states also increases. While in a single period model this does not cause much computational burden, it suggests that it will lead to a state explosion in the multi-period case. In the future we plan to investigate the approach of finding the state that causes the arbitrage opportunity and eliminating/modifying it rather than just adding more states.

Table 1. States added to guarantee arbitrage free tree

Assets	States Added	Time (secs)
2	4	0
5	14	0
10	27	0.01
15	77	0.14
20	170	0.79

Table 2. Solution Times Moment Problem

Assets	Time (secs)	Variables
2	820	105
5	1230	111
10	4160	124
15	23316	184
20	52544	242

Table 3. Solution Times Portfolio Selection

Assets	M-V	M-V-S	M-V-K	M-V-S-K
2	0.01	56	36	34
5	0.3	138	101	204
10	1.3	250	195	195
15	4	375	433	519
20	14.8	551	776	762

In table 2 we show the time required to solve the moment problem. The number of variables differ from one run to the next. This is because the number of states that are needed to guarantee the arbitrage free property differ from run to run (since they are randomly generated). In all runs we added fifty more (non-constant) states. The resulting problem given by (2) was then solved using the algorithm described above. The times shown are the averages for ten runs for problems with 2, 5, and 10 assets. Due to the large amount of time required to solve the larger problems (15 and 20 assets) the times reported are from a single run. Table 3 details the time required to generate a point on the efficient frontier for the four versions of the portfolio selection problem we considered in this paper. The first is the mean-variance (M-V) model, this is obtained by setting $\gamma = \delta = 0$ (we used this model to test the quality of the solutions provided by the algorithm). Similarly M-V-S, M-V-K and M-V-S-K stand for Mean-Variance-Skewness, Mean-Variance-Kurtosis and Mean-Variance-Skewness-Kurtosis respectively. We realize that providing the computation times is not the best way to judge the speed of an algorithm. However, one of the aims of this paper is to show how one can use a stochastic global optimization algorithm to solve a financial planning problem. Even though there are many open questions, and some of our assumptions may be too stringent, we believe that the computation times tabulated below show the feasibility of this approach. In figures 1 and 2 we show some 3-dimensional efficient frontiers for the M-V-S and M-V-K problems respectively. The gaps that appear in the frontiers are due to the way we generate the frontier. If we used constraints, instead of weights, to express the preferences of the DM, then we believe that the frontier would look more smooth. However, a formulation using constraints would lead to an optimization problem that could not be solved by our global optimization solver. We plan to address this deficiency in the future. Figures 3 and 4 show efficient frontiers using the M-V-S-K model. In figure 3 we plot the first three measures of interest, while in figure 4 we plot the mean, variance and kurtosis of the portfolio.

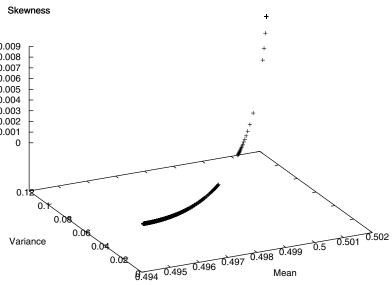

Fig. 1. Mean-Variance-Skewness

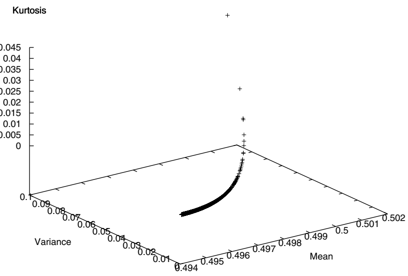

Fig. 2. Mean-Variance-Kurtosis

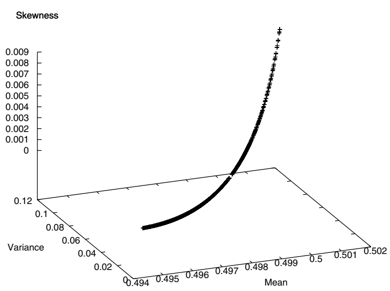

Fig. 3. Mean-Variance-Skewness-(Kurtosis)

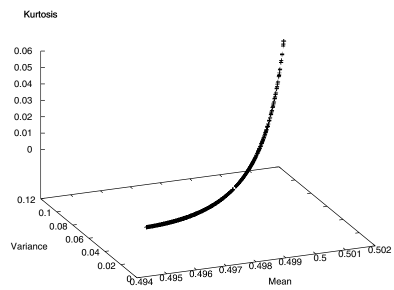

Fig. 4. Mean-Variance-(Skewness)-Kurtosis

6 Conclusions

We considered the computational challenges associated with the global optimization of a financial planning model. We proposed a simple extension to the classical Markowitz model; in the proposed model higher order moments were included using scalar weights. The scenario generation problem was also addressed by matching the first four central moments of the postulated distribution. We also addressed the issue of imposing the arbitrage free property to the generated scenario tree. A stochastic algorithm was proposed for the two models. Our initial numerical results show that problems of realistic size can be solved using the proposed framework.

References

1. Markowitz, H.M.: Portfolio selection. J. Finance **7** (1952) 77–91
2. Markowitz, H.M.: The utility of wealth. J. Polit. Econom. (1952) 151–158
3. Dupacova, J., Consigli, G., Wallace, S.: Scenarios for multistage stochastic programs. Ann. Oper. Res. **100** (2000) 25–53 (2001)
4. Gülpınar, N., Rustem, B., Settergren, R.: Simulation and optimization approaches to scenario tree generation. J. Econom. Dynam. Control **28** (2004) 1291–1315
5. Prékopa, A.: Stochastic programming. Volume 324 of Mathematics and its Applications. Kluwer Academic Publishers Group, Dordrecht (1995)
6. Parpas, P., Rustem, B., Pistikopoulos, E.N.: Linearly constrained global optimization and stochastic differential equations. Accepted in the J. Global Optim. (2006)
7. Aluffi-Pentini, F., Parisi, V., Zirilli, F.: Global optimization and stochastic differential equations. J. Optim. Theory Appl. **47** (1985) 1–16
8. Chiang, T., Hwang, C., Sheu, S.: Diffusion for global optimization in $\mathbf{R}^n$. SIAM J. Control Optim. **25** (1987) 737–753
9. Geman, S., Hwang, C.: Diffusions for global optimization. SIAM J. Control Optim. **24** (1986) 1031–1043
10. Gidas, B.: The Langevin equation as a global minimization algorithm. In: Disordered systems and biological organization (Les Houches, 1985). Volume 20 of NATO Adv. Sci. Inst. Ser. F Comput. Systems Sci. Springer, Berlin (1986) 321–326
11. Konno, H.: Applications of global optimization to portfolio analysis. In C., A., P., H., G., S., eds.: Essays and Surveys in Global Optimization. Springer (2005) 195–210
12. Maranas, C.D., Androulakis, I.P., Floudas, C.A., Berger, A.J., Mulvey, J.M.: Solving long-term financial planning problems via global optimization. J. Econom. Dynam. Control **21** (1997) 1405–1425 Computational financial modeling.
13. Steinbach, M.C.: Markowitz revisited: mean-variance models in financial portfolio analysis. SIAM Rev. **43** (2001) 31–85 (electronic)
14. Parpas, P., Rustem, B.: Entropic regularization of the moment problem. Submitted (2005)
15. Bertsimas, D., Sethuraman, J.: Moment problems and semidefinite optimization. In: Handbook of semidefinite programming. Volume 27 of Internat. Ser. Oper. Res. Management Sci. Kluwer Acad. Publ., Boston, MA (2000) 469–509
16. Klaassen, P.: Discretized reality and spurious profits in stochastic programming models for asset/liability management. E.J of Op. Res. **101** (1997) 374–392
17. Harrison, J., D.M., K.: Martingales and arbitrage in multiperiod securities markets. J. Econom. Theory **20** (1979) 381–408
18. Zirilli, F.: The use of ordinary differential equations in the solution of nonlinear systems of equations. In: Nonlinear optimization, 1981 (Cambridge, 1981). NATO Conf. Ser. II: Systems Sci. Academic Press, London (1982) 39–46
19. Recchioni, M.C., Scoccia, A.: A stochastic algorithm for constrained global optimization. J. Global Optim. **16** (2000) 257–270