J Optim Theory Appl (2012) 155:707–722
DOI 10.1007/s10957-012-0074-x

Risk-Return Trade-off with the Scenario Approach in Practice: A Case Study in Portfolio Selection

B.K. Pagnoncelli · D. Reich · M.C. Campi

Received: 6 April 2011 / Accepted: 28 April 2012 / Published online: 12 May 2012

Abstract We consider the scenario approach for chance constrained programming problems. Building on existing theoretical results, effective and readily applicable methodologies to achieve suitable risk-return trade-offs are developed in this paper. Unlike other approaches, that require solving non-convex optimization problems, our methodology consists of solving multiple convex optimization problems obtained by sampling and removing some of the constraints. More specifically, two constraint removal schemes are introduced, one greedy and the other randomized, and a comparison between them is provided in a detailed computational study in portfolio selection. Other practical aspects of the procedures are also discussed. The removal schemes proposed in this paper are generalizable to a wide range of practical problems.

Keywords Chance constrained programming · Scenario approximation · Portfolio selection

Communicated by Johannes O. Royset.

This research was funded by ANILLO Grant ACT-88 and Basal project CMM, Universidad de Chile.

B.K. Pagnoncelli (✉) · D. Reich
Escuela de Negocios, Universidad Adolfo Ibañez, Diagonal las Torres 2640 oficina 533-C, Peñalolén, Santiago, Chile
e-mail: bernardo.pagnoncelli@uai.cl

D. Reich
Ford Research and Advanced Engineering, Dearborn, MI 48124, USA

M.C. Campi
Department of Information Engineering, University of Brescia, Via Branze 38, 25123 Brescia, Italy

1 Introduction

Scenario-based (or sample-based) approximations to chance constrained programming (CCP) problems has been an active area of research in recent years and several important theoretical developments have emerged, leading to different strategies to solve those problems. In [1], the authors employ a large deviation-type approximation, referred to as the "Bernstein approximation," which is convex and can be efficiently solved. Another approach, described in [2, 3], requires finding the finite set of p-efficient points of a probability distribution in order to derive equivalent problem formulations. In the case where uncertainty is given by a discretely distributed random vector, it is possible to formulate an equivalent mixed integer program of large dimension, which includes "big-M" constraints. Such problems are computationally challenging, and some promising lines of research have appeared recently in the literature. In [4], the authors propose a new method based on irreducibly infeasible subsystems and show some promising computational experiments in a vaccine allocation problem. The work of [5] and the more recent publication [6] propose an efficient method to solve chance constrained problems, in which only the right-hand side is random and this random vector follows a finite distribution. In [7], the author takes a different path and proposes a pattern recognition method to efficiently solve stochastic programs with a larger number of scenarios.

This paper presents a case study in portfolio selection based on the scenario approach discussed in [8–10]. Scenario-based approaches do not require restrictive assumptions on the distribution of the random parameters, which increases their generality and flexibility. For instance, dependence among random parameters can be incorporated into sampling frameworks avoiding the intractable task of computing multidimensional integrals. If the dependence is given, e.g., by a known copula, generating samples requires minimal computational effort [11]. While some computer languages require customized header files in order to be able to generate random numbers, other software packages for scientific computing include built-in routines for random number generation, making the implementation of these approaches quite straightforward.

On the other hand, the challenge in applying scenario-based frameworks is that they tend to produce conservative solutions for their corresponding CCP problems. In [12], some of these methods were tested in a portfolio selection problem and the authors concluded that the methods indeed proved to be too conservative. In this paper, we explore the idea, recently introduced in [10], of utilizing *constraint removal* as a means for achieving high-quality, nonconservative solutions to CCP problems, an idea that presents practical challenges, mainly tractability.

More precisely, while [10] has set the theoretical basis for a scenario approach with constraint removal, the main goal of this paper is to provide guidelines for using this framework in practice in a portfolio selection problem. To this purpose, we further develop the framework's theoretical foundation, and moreover provide simulation studies to confirm our analysis. In [10], Campi and Garatti state that any constraint removal algorithm can be employed and provide a theoretical bound that relates constraint removal to feasibility. We specialize their work by proposing two distinct constraint removal schemes: one greedy and the other randomized.

The important consequence of removing constraints in the context described in [10] is that the found solution is less conservative, and thereby this approach overcomes the difficulties pointed out in [12] and allows a practical application of the scenario methodology to CCP problems. While our numerical study is conducted on a portfolio selection problem, the methodology of constraint removal and the findings of this paper have broader scope and can be readily applied to other chance constrained problems.

The remainder of the paper is organized as follows. In Sect. 2, we define the portfolio selection problem on which we later apply our methodology. The solution methodologies for constraint removal are introduced in Sect. 3, and Sect. 4 presents an extensive computational study in portfolio selection. Section 5 summarizes our contributions and discusses future research directions.

2 A Portfolio Selection Problem

We consider a portfolio problem taken from [12]. A robust version of this problem was considered in [13]. The goal of this problem is to allocate a percentage x of an initial capital amongst n possible assets, so that the expected growth of the capital is maximized while requiring some minimum desired return v per unit of capital invested. In order to increase the potential for growth, the investor is willing to accept a certain degree of risk, and the minimum desired return v is only enforced with a high probability $1-\varepsilon$, where ε represents the risk. For example, if $v=0.90$ and $\varepsilon = 5\ \%$, the investor wishes to have less than a $5\ \%$ chance of losing more than 0.10 of the initial investment. This portfolio problem can be formulated as the following chance constrained program:

$$\operatorname{Max} \bar{\xi}^T x \qquad \text{s.t.} \quad \mathsf{P}\{\xi^T x \geq v\} \geq 1-\varepsilon, \quad \sum_{i=1}^{n+1} x_i = 1, \quad x \in [0,1]^{n+1}, \tag{1}$$

where ξ is a random vector representing the returns of the $n+1$ possible assets[1] and $\bar{\xi} = \mathbb{E}[\xi]$ is its expected value. Each component of ξ represents the quotient between the price at some future time, the unknown, and the current price, which is known. We assume that x_{n+1} represents the percentage invested in a risk-free asset with a constant return of 1, so that a percentage of the initial capital can be simply retained as cash without generating additional returns. The inclusion of this risk-free asset is important from a practical perspective and, moreover, it ensures that, for all $v \in [0,1]$, there is always a feasible investment strategy.

In our case study, data are generated artificially from a known distribution, which we estimated from historical data. In practice, when applying our methodology, as explained later, one could work directly with the historical data. However, the purpose of this paper is to introduce and analyze constraint removal methodologies, and using

[1] Note that the dimension of our problem is actually n, since once percentages $x_1, \ldots, x_n$ are allocated, x_{n+1} must become the remaining unallocated percentage.

an artificial distribution turns out to be convenient because one can a posteriori test the quality of the obtained results, since the underlying distribution is known.

We assume that ξ follows a multivariate lognormal distribution, which is consistent with the classical Black and Scholes framework. It is also an interesting experimental assumption, since the chance constraint in (1) becomes a sum of lognormals, which does not have a known distribution and, therefore, cannot be converted into a deterministic constraint. In contrast, under assumptions of normality, it is well known (e.g., [14], page 16) that the chance constraint in (1) can be converted via the error function into a deterministic second-order conic constraint.

This portfolio model is intended to serve as an example to test the scenario approach methodologies for CCPs. In terms of risk measures, the chance constraint in (1) is equivalent to a Value-at-Risk constraint [15]. However, from an applied financial perspective, the portfolio model we consider is naive in that it does not incorporate several crucial features of real markets, e.g., cost of transactions, short sales, lower and upper bounds on holdings, etc. For more realistic models, we refer the reader, e.g., to [16], where the authors include market frictions and discuss the best distribution function for asset returns, and to [17], where the author incorporates lower and upper bounds on the holdings of both individual assets and asset groupings.

In the next section, we contextualize problem (1) within the more general CCP framework and define concepts that we use in developing our scenario-based solution methodologies.

3 Chance Constrained Problems

Problem (1) is a particular case of the general chance constrained problem:

$$\operatorname*{Max}_{x\in X} c^T x \qquad \text{s.t.} \quad \mathsf{P}\big\{g(x,\xi)\le 0\big\}\ge 1-\varepsilon, \tag{2}$$

where c is a deterministic vector in $\mathbb{R}^n$, ξ is the random vector, $1-\varepsilon\in[0,1]$ is the *reliability level* and $X\in\mathbb{R}^n$ is a convex and closed set. We assume the convexity of $g(\cdot,\xi)$ as a function of x, while no assumption on ξ is made. Multiple constraints $g_j(x,\xi)\le 0$, $j=1,\dots,m$, can be equivalently replaced via the max-function, so, without any loss of generality, we can consider the single constraint case. Of course, in some cases (e.g., random technology matrix), the max operator may hide some useful property of the original functions, and the difficulty of the resulting problem may increase significantly. In the context of problem (1), $g(x,\xi)=v-\xi^T x$. Linearity of the objective function is without any loss of generality, since any problem with a convex objective function can be written with a linear objective plus a convex constraint by an epigraphic reformulation. For instance, if the objective function is to minimize $f(x)$, then, by introducing an auxiliary variable θ, one can consider the equivalent problem of minimizing θ subject to having θ greater or equal to $f(x)$, which has a linear objective and a convex constraint.

The following definition refers to the complement of the left-hand side of the chance constraint, which is the crucial information for determining if a candidate solution satisfies the required reliability level.

Definition 3.1 Let $x \in X$ be a feasible solution for (2). The violation probability of x is defined by

$$V(x) := \mathsf{P}\{\xi : g(x,\xi) > 0\},$$

where we assume that $\{\xi : g(x,\xi) > 0\}$ is measurable.

In the context of (1), the violation probability $V(x)$ measures the probability with which the minimum required level of return v is not achieved.

3.1 The Scenario Approach

Sample approaches have been studied extensively over the last 10 years as a means of identifying good candidate solutions to problem (2). They have been referred to as randomized or scenario approaches in [1, 8, 9, 18], and as sample average approximation in [5, 12, 19]. The idea behind these approaches is straightforward and intuitive. Aside from trivial cases, chance constrained problems cannot be converted into tractable deterministic optimization problems. Therefore, in order to solve these problems, we have to rely on approximations. The following definition provides one such possibility.

Definition 3.2 Let $\xi^1, \dots, \xi^N$ be an independent and identically distributed sample of size N, extracted according to probability P (this is the same probability appearing in the definition of the chance constraint in (2)). The *scenario program* associated with problem (2) is defined as

$$\operatorname*{Max}_{x \in X} c^T x \qquad \text{s.t.} \quad g\left(x, \xi^i\right) \le 0, \quad i = 1, \dots, N. \tag{3}$$

Calafiore and Campi [8] and Campi and Garatti [9] have shown how the sample size N should be chosen in problem (3) to ensure that, with an arbitrarily high confidence, the optimal solution $\hat{x}_N$ of (3) is feasible for the original problem (2). However, there is no guarantee that the solution $\hat{x}_N$ of (3) will obtain a value $c^T \hat{x}_N$ close to the true optimal value $c^T x^*$ of (2).

In fact, we have reasons to expect that the solution $\hat{x}_N$ will be conservative for the original CCP. Indeed, based on results derived in [9], the distribution function of $V(\hat{x}_N)$ is bounded by a beta distribution with parameters n and $N - n + 1$, and imposing that $V(\hat{x}_N) \le \varepsilon$ holds with high confidence implies that $V(\hat{x}_N)$ will be much less than ε in many cases, resulting in a conservative solution.

3.2 The Scenario Approach with Constraint Removal

To identify less conservative solutions $\hat{x}_N$, we build upon a general framework introduced by Campi and Garatti in [10] for relaxing problem (3). Their approach allows one to remove k constraints out of the N scenario constraints in (3); as shown in [10], feasibility for problem (2) is retained with confidence $1 - \beta$, provided the following inequality is satisfied:

$$\binom{k+n-1}{k} \sum_{j=0}^{k+n-1} \binom{N}{j} \varepsilon^j (1-\varepsilon)^{N-j} \le \beta, \tag{4}$$

($\binom{m}{p}$ is the binomial coefficient: $\binom{m}{p} = \frac{m!}{p!(m-p)!}$). This result holds true independently of probability P, that is, (4) is a distribution-free result. To achieve the least conservative solution possible, for a given N, ε, n and β, we always choose, in the following, the largest k satisfying (4). We would like to stress that we are not removing constraints from the original problem (2), but from the scenario approximation (3). Therefore, the well-known theory of sensitivity analysis with respect to the constraints' coefficients and KKT conditions are not readily applicable here because the constraints removed do not change the original feasible set of problem (2).

The result in [10] holds true independently of the procedure for eliminating k constraints. On the other hand, the removed constraints impact the value of $c^T \hat{x}_N$. One optimal way of removing constraints consists in discarding those constraints that lead to the largest possible improvement of the cost function. This approach is implemented by the following integer program, which has been considered in [12] and [19]:

$$\begin{aligned} &\operatorname*{Max}_{x\in X} c^T x \qquad \text{s.t.} \quad g\bigl(x,\xi^i\bigr) - MZ_i \le 0, \quad i = 1,\dots,N, \\ &\sum_{i=1}^{N} Z_i \le k, \quad Z \in \{0,1\}^N, \end{aligned} \tag{5}$$

where M is a constant large enough so that, if $Z_i = 1$, then the constraint is satisfied for any candidate solution x. Formulation (5) is a relaxation of (3) since the allowed violation expands the feasible set of (3). For $k = 0$, these two formulations are equivalent.

By construction, problem (5) provides a framework for optimally selecting the constraints to be removed based on Campi and Garatti's inequality (4). However, solving (5) may be computationally challenging due to the increase in complexity from (3) to (5) that arises from the introduction of one binary variable per each of the N scenarios.

3.3 Greedy and Randomized Constraint Removal Procedures

In this paper, we introduce and study greedy and randomized constraint removal methods for obtaining feasible solutions to problem (5), by solving multiple problems of the form (3). Our greedy procedure iteratively removes k constraints. At each iteration i, an initial linear program of form (3) with $N - i + 1$ constraints is solved to determine the set of n_i active constraints. We then remove the active constraints one at a time and solve the corresponding n_i additional linear programs of form (3), each with $N - i$ constraints. The final step at each iteration is to identify the active constraint whose elimination yields the greatest improvement in the objective value, and remove it. The greedy removal algorithm requires solving $1 + \sum_{i=1}^{k} n_i$ problems of form (3).

We also introduce a randomized removal algorithm, where at each iteration we at random remove one active constraint, so that they are all equally likely to be removed. This improves tractability by requiring that only $1 + k$ problems of the form (3) are solved.

4 Computational Study

We designed computational experiments to study the tractability of constraint removal approaches for obtaining high-quality solutions to the original portfolio optimization chance constrained problem. While all our experiments were performed in the context of a portfolio optimization problem, the constraint removal scheme proposed here is directly and readily applicable to other CCP problems.

The key points we present in this section can be summarized as follows:

1. In the limit as $N \to \infty$, k/N approaches ε in inequality (4) (this result has been proven in [10]). We provide examples that demonstrate the practical limitations of this result (4) due to computational issues.
2. The practical limitations in the use of inequality (4) can be alleviated by the use of efficient constraint removal procedures, specifically of greedy and randomized types.
3. We compare the greedy and the randomized procedures and show that the latter performs nearly as well as the former for sufficiently large N. However, the latter has the advantage of being much less computationally demanding.
4. Importantly, one can run the scenario algorithm more times and pick the maximum solution obtained over the different runs. We show that this is an effective technique to improve the solution, while rigorously preserving high confidence that the violation is below ε.
5. Upper bounds for the true optimal value, which can be efficiently obtained, can be used to provide optimality gaps with arbitrarily high probability.

4.1 Selecting the Sample Size and the Number of Constraints to Be Removed

Intuitively, it seems advantageous to choose the number k of constraints to remove, and the number of scenarios N in (4), so that k/N becomes as close as possible to the original violation ε. Theoretically, as $N \to \infty$, we can choose k/N arbitrarily close to ε, while guaranteeing feasibility to problem (2). However, choosing excessively large values of N would result in computationally intractable problems.

For our asset portfolio problem (1) with $n = 20$ and for $\beta = 10^{-9}$, Fig. 1 shows the convergence of k/N to $\varepsilon = 5\ \%$ as N increases. The x-axis is in a log base 2 scale and, for example, over one million scenarios would be needed to achieve values of k/N greater than 4.5 %. In our experiments, we selected values of N up to 20000, with ratios of k/N up to 2.9 %. (Recall: This does not imply that the solution will have a violation level of 2.9 %; rather, the solution is guaranteed with confidence greater than or equal to $1 - \beta = 1 - 10^{-9}$ to have a violation level of less than $\varepsilon = 5\ \%$).

4.2 Design of the Experiments

We consider three minimum desired return values v: 0.85, 0.90, and 0.95. For each one, we consider the four sample sizes N ranging from 2500 up to 20000 listed in Table 1, along with the number k of removed constraints, ratio k/N and the confidence parameter β. The value of β was obtained from inequality (4). It is the probability of

Fig. 1 Number of scenarios versus k/N, $\beta = 10^{-9}$

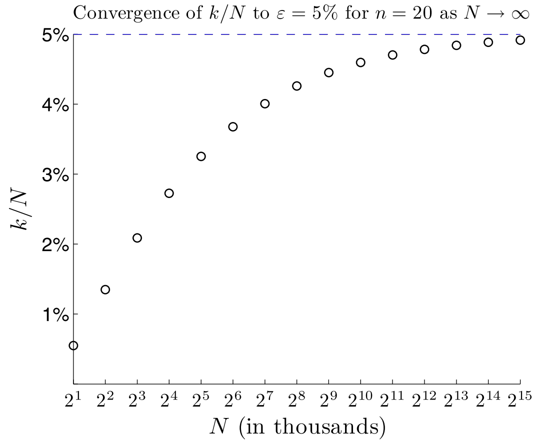

Table 1 Sample sizes N, along with the corresponding max number of constraint removals k, ratio k/N, and probability β of obtaining an invalid solution to the original chance constrained problem with $\varepsilon = 5\ \%$

N	k	k/N	β
2500	18	0.007	7.16×10^{-11}
5000	76	0.015	9.67×10^{-11}
10000	220	0.022	1.57×10^{-12}
20000	582	0.029	9.93×10^{-9}

obtaining an infeasible solution to the original chance constrained problem (1) with $\varepsilon = 5\ \%$.

Our experiments consisted of 30 runs for each combination of the minimum desired return v and N. Different seeds were used in different runs, and for each run we used the same seed for both the greedy and randomized constraint removal procedures, so that the performance of these procedures can be compared objectively.

Designing our experiment with 30 runs for each set of parameters serves two purposes. The first is that it allows us to gain experimental evidence on the variability of a highly complex and analytically evasive optimization method. The second benefit is of a methodological nature: we can select the best solution over a group of 30 runs (instead of a single run) as a means to achieve a solution carrying an improved performance. Regarding this latter point, the following theoretical observation is in order: a repeated application of the result stated in (4) allows one to perform several runs, take the best solution and still be guaranteed that the violation is no more than ε with large confidence. Indeed, we have that

$$\mathsf{P}\Big\{\max_{r=1,\dots,30} V\big(\hat{x}_N^r\big) \le \varepsilon\Big\} = \mathsf{P}\big\{V\big(\hat{x}_N^1\big) \le \varepsilon, \dots, V\big(\hat{x}_N^{30}\big) \le \varepsilon\big\} \ge 1 - 30\beta,$$

where r is a parameter running over the various runs. Since β is very small, we can guarantee that the best solution over 30 runs has a probability of violation below ε with high confidence; this relates to the fact explained in [10] that confidence is very "cheap," so that having 30β instead of β impacts only very marginally the values of k and N.

Table 2 The 20 assets included in our numerical experiments (with their ticker symbols in parentheses)

Assets	
Boeing Company (BA)	Chevron Corporation (CVX)
Walt Disney Company (DIS)	Eastman Kodak Company (EK)
Ford Motor Company (F)	Goodyear Tire & Rubber Company (GT)
Honeywell International (HON)	Johnson & Johnson (JNJ)
Kellogg Company (K)	Coca-Cola Company (KO)
Kroger Company (KR)	Lockheed Martin Corporation (LMT)
McDonald's Corporation (MCD)	3M Company (MMM)
Merck & Company Inc (MRK)	Pepsico Inc (PEP)
Procter & Gamble Company (PG)	Wal-Mart Stores Inc. (WMT)
Exxon Mobil Corporation (XOM)	Xerox Corporation (XRX)

The final component of our experiment consists of computing upper bounds for the values of the original problem (1) following [1] and [12]. These upper bounds are obtained by solving many small sampled problems without constraint removal and using the Uth largest objective value, where the choice of U comes from a binomial-type formula (see p. 7 of [12] for more details). In this way, we obtain interesting information on the optimality gap of the solution obtained through the scenario approach.

We implemented and solved our model using Mosel Xpress. All computations were performed on identical machines each having x86_64 architecture, an Intel Xeon @2.5 GHz processor, 16 GB RAM, and running CentOS 5 Linux.

4.3 Data

We gathered data from Yahoo Finance[2] for the 20 assets in Table 2. For all assets, we used monthly data covering a 25 year span from the beginning of 1985 through the beginning of 2010. Additionally, we include a risk-free asset with a return equal to 1, that is, any amount allocated to the risk-free asset would be held as cash.

To smooth the data, we considered the price of each month over the closing price one year prior. We assumed these quotients follow a multivariate lognormal distribution and estimated the parameters of the distribution using the usual unbiased estimators for means and covariances [20].

4.4 Computational Results: Greedy Constraint Removal

Table 3 provides both the mean and maximum objective function values, as well as the mean percentage invested in the risk-free asset, over 30 runs for each set of parameters. In addition, Table 3 provides upper bounds for the optimal values of the chance constrained problem (1) for each minimum desired return considered obtained, as indicated at the end of Sect. 4.2. As explained in Sect. 4.2, one can select the allocation

[2] http://finance.yahoo.com/.

Table 3 The conservativeness of solutions decreases as k/N increases, where the mean and max are taken over 30 runs that use the greedy removal algorithm

k/N	Mean objective function	Mean risk-free asset	Max objective function
$v = 0.85$			
0.007	1.0370	0.34	1.04071
0.015	1.0406	0.28	1.04231
0.022	1.0434	0.23	1.04505
0.029	1.0459	0.19	1.04709
Upper bound = 1.06497			
$v = 0.90$			
0.007	1.0248	0.56	1.02666
0.015	1.0272	0.51	1.02807
0.022	1.0289	0.49	1.02946
0.029	1.0307	0.46	1.03131
Upper bound = 1.05208			
$v = 0.95$			
0.007	1.0124	0.78	1.01306
0.015	1.0136	0.76	1.01409
0.022	1.0143	0.74	1.01486
0.029	1.0153	0.73	1.01582
Upper bound = 1.02635			

leading to the maximum objective value while preserving high confidence that the violation is below 5 %.

There are two key observations, which hold true for all three levels of 0.85, 0.90, and 0.95 minimum desired returns. The first is that, as k/N increases, the solutions obtained become less and less conservative. This decrease in conservativeness is evidenced both in the mean and maximum objective values, as well as in the decreasing percentage of capital invested in the risk-free asset. The second important observation is that the solutions achieved for larger sample sizes (which correspond to larger k/N values) show less variation, as it can be seen by inspecting the gap between the mean and maximum statistics is shrinking as N increases. In addition, we observe that the impact of the value k/N on the percentage allocated to the risk-free asset is larger for lower values of v. This indicates that, for low values of v, the scenario approach can introduce conservativeness if k/N is not selected large enough.

Since no closed-form exists for evaluating the violation $V(\cdot)$ from Definition 3.1, we computed estimates for the optimal solutions via Monte Carlo sampling with a required sample standard deviation of 0.001. Table 4 provides the minimum, mean, and maximum violation estimates $\hat{V}(\cdot)$ over the 30 runs. It can be noted that the violation is always below the required risk $\varepsilon = 5$ %; indeed, the fact that the violation is upper bounded by 5 % is a result theoretically guaranteed with the high confidence $1 - \beta$ given in Table 1.

In Table 4, the real violation is well below 5 %. This is due to the stochastic nature of the violation; specifically, the violation is subject to stochastic fluctuation (com-

Table 4 Violation estimate statistics for 30 runs

k/N	Min violation	Mean violation	Max violation
$v = 0.85$			
0.007	0.8 %	1.2 %	2.0 %
0.015	1.3 %	1.8 %	2.1 %
0.022	1.9 %	2.4 %	2.8 %
0.029	2.6 %	2.9 %	3.2 %
$v = 0.90$			
0.007	0.7 %	1.2 %	1.8 %
0.015	1.4 %	1.9 %	2.4 %
0.022	2.1 %	2.4 %	2.7 %
0.029	2.6 %	3.0 %	3.3 %
$v = 0.95$			
0.007	0.8 %	1.3 %	1.8 %
0.015	1.5 %	1.9 %	2.2 %
0.022	1.9 %	2.3 %	2.8 %
0.029	2.5 %	3.0 %	3.4 %

pare, e.g., the minimum with the maximum violations in Table 4), so that the practical implication of ensuring that the violation stays below 5 % with high confidence is that the achieved violation is normally not close to the limit ε.

4.5 Computational Results: Greedy vs. Randomized Constraint Removal

Figures 2 and 3 compare the performance of the greedy and randomized constraint removal algorithms. In Fig. 2, the comparison is for a single run, whereas Fig. 3 compares the mean performance over 30 runs. One salient feature that is worth noting is that, while for smaller values of k the greedy algorithm (GA) outperformed the randomized algorithm (RA), for larger values of k the gap in terms of objective function value became almost negligible. An explanation of this behavior can be found in Table 5.

Table 5 shows that the mean percentage of constraints removed both by the GA and the RA (what we call overlap) over 30 runs remains roughly constant for a given number of removals k while varying N, while it increases significantly with k. Indeed, a striking statistic in Table 5 is that the mean percentage overlap increases significantly as the number of removals increases from 18 to 582, climbing from just over half to 92 %. Another way of interpreting the results is that, for values of k much larger than $n = 20$ (the number of optimization variables), such as 220 and 582, a constraint removed by the GA will eventually be removed by the RA as well with high probability. This explains why in Fig. 2, for $N = 20000$ and $k = 582$, both curves are almost indistinguishable: the GA and the RA removed essentially the same set of constraints. As a practical guideline, after k and N have been selected, we suggest using RA if k is large enough, as compared with n due to its superior running time, which is shown in Table 6.

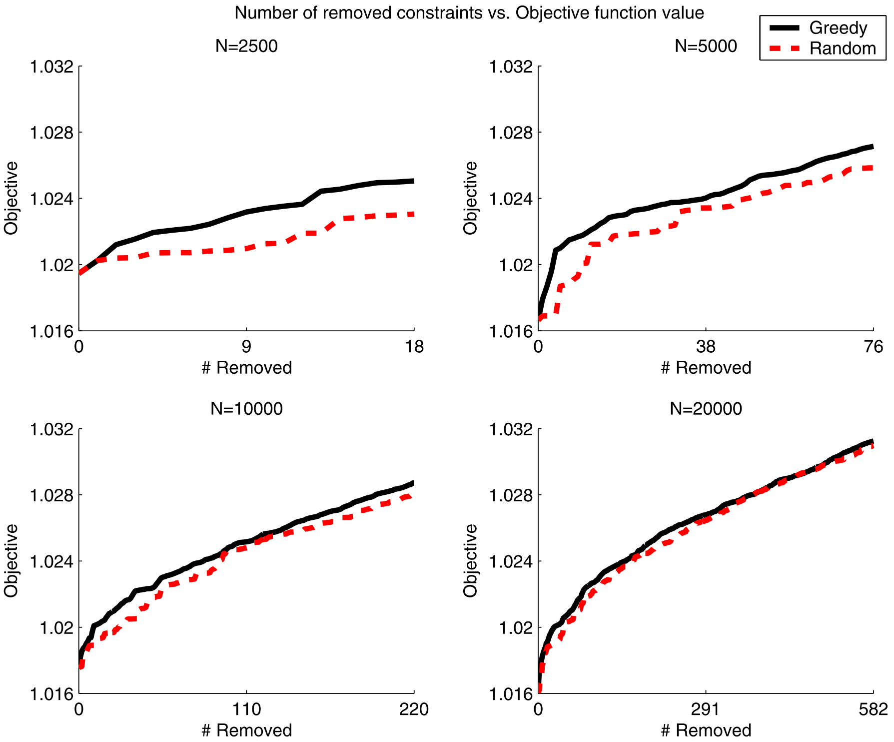

Fig. 2 Single run

4.6 Large Scale Experiments

We ran a second set of experiments with 200 assets to demonstrate the scalability of the RA approach and confirm its suitability for solving large scale portfolio problems with a VaR constraint. We obtained monthly data from Yahoo Finance spanning the 20-year period between 1990 and 2010. Our goal with the low-dimensional 20 asset experiments was both to investigate the relationship between k/n and ε and to compare the random and the greedy removal approaches. We established that, for larger sample sizes, the two approaches become almost indistinguishable, with RA taking significantly less computational time.

The experiments with 200 assets confirm these findings. By running both GA and RA with 20000 samples and 582 removals, we can see in Table 7 that for $v = 0.85, 0.90$ and 0.95 the objective function values obtained for both methods are nearly identical. This provides further evidence that RA should be used instead of GA, because its computational time is significantly lower.

It is interesting to compare Table 7 with Table 3. By starting with the same number of samples (20000) and removals (582), as in our original 20 asset experiment, we observe that the additional investment opportunities significantly increase the expected returns. Moreover, with 200 assets available, no investment in the risk free asset was needed for $v = 0.85$ and $v = 0.90$. Even for $v = 0.95$, the amount invested

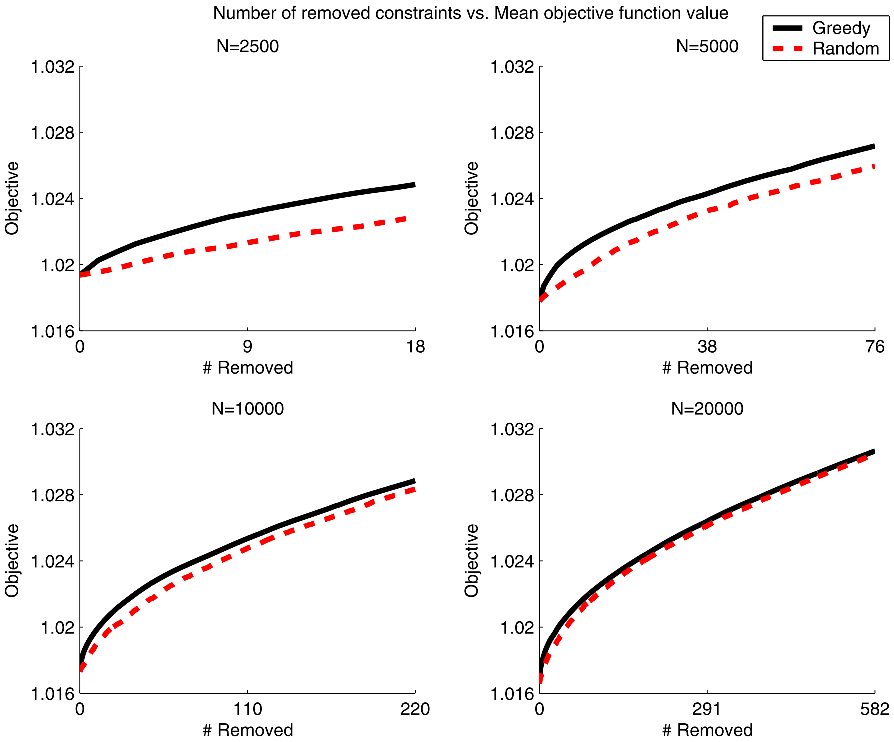

Fig. 3 Averages

in the risk free asset is still only 36.7 % (RA) and 34.9 % (GA), which is roughly half of what is invested in the 20 assets case. On the other hand, the guaranteed risk with 200 assets is lower than with 20 assets, and an application of formula (4) shows that with $\beta = 9.93 \times 10^{-9}$ (as in Table 1 for the 20 assets experiment) one obtains $\varepsilon = 9.5$ %. Thus, with a number of assets 10 times larger, ε does increase, but only by a factor 2 or so. This relatively moderate increase of ε is related to constraint removal, a fact that can also be appreciated by looking at formula (4), where n appears always summed to k so that k smooths down the effect of large variations in n.

We ran additional experiments using only RA for sample sizes 40000 and 80000 with 1164 and 2328 removals, respectively, maintaining a constant k/N ratio of roughly 2.9 %. Table 8 provides the obtained objective function values. Comparing Table 8 with Table 7, we see that the objective values are alike in the two cases. Taking $\beta = 9.93 \times 10^{-9}$ in (4), with $N = 40000$ gave $\varepsilon = 7.4$ %, while $\varepsilon = 3.8$ % was obtained for $N = 80000$.

With 20 assets and $N = 20000$, the mean execution time for GA was more than 10 times greater than for RA, as shown in Table 6. With 200 assets and $N = 20000$, the ratio increased to 16.5 and the execution time for RA is over an hour, suggesting that GA may become impractical for larger problems. Table 9 provides detailed running times for RA with 200 assets.

Table 5 Mean percentage overlap of constraints removed

N	Mean overlap after first 18 removals	Mean overlap after first 76 removals	Mean overlap after first 220 removals	Mean overlap after first 582 removals
$v = 0.85$				
2500	55 %			
5000	56 %	79 %		
10000	58 %	79 %	88 %	
20000	55 %	77 %	87 %	92 %
$v = 0.90$				
2500	54 %			
5000	60 %	78 %		
10000	59 %	78 %	87 %	
20000	58 %	78 %	87 %	92 %
$v = 0.95$				
2500	58 %			
5000	56 %	78 %		
10000	60 %	80 %	88 %	
20000	55 %	79 %	87 %	92 %

Table 6 Average running times for the greedy and the randomized algorithms over 30 runs

N	k	Mean running time greedy (s)	Mean running time randomized (s)
2500	18	15.1	2.3
5000	76	138.0	14.8
10000	220	875.3	86.0
20000	582	5412.4	504.6

Table 7 Results for $N = 20000$, $k = 582$ and 200 assets

Algorithm	Mean objective function	Mean risk-free asset	Max objective function
$v = 0.85$			
RA	1.16052	0.0	1.16133
GA	1.16112	0.0	1.16197
$v = 0.90$			
RA	1.13903	0.0	1.14161
GA	1.14050	0.0	1.14249
$v = 0.95$			
RA	1.07812	0.37	1.08048
GA	1.08042	0.35	1.08311

Table 8 Results for $N = 40000$ and 80000 samples and $k = 1164$ and 2328, respectively

Sample size	Mean objective function	Mean risk-free asset	Max objective function
$v = 0.85$			
40000	1.16071	0.0	1.16154
80000	1.16096	0.0	1.16130
$v = 0.90$			
40000	1.13964	0.0	1.14093
80000	1.140129	0.0	1.14099
$v = 0.95$			
40000	1.07878	0.37	1.08180
80000	1.07864	0.37	1.07995

Table 9 Average running times with 200 assets for RA over 30 runs

N	k	Mean running time randomized (s)
20000	582	5521
40000	1164	26307
80000	2328	120535

5 Summary and Conclusions

Scenario programs with constraint removal have been considered in this paper, and practical constraint removal methods have been proposed. Specifically, two easily implementable polynomial-time methods have been discussed, one greedy and the other randomized. Our methodology has been applied to a value-at-risk portfolio problem, in which an investor wants to maximize the expected return while controlling the probability of significant losses. We ran two sets of experiments, the first with 20 assets aiming at testing the methodology and obtaining insight about the removal schemes, and a second one with 200 assets to test the feasibility of our approach for problems with a larger pool of assets. It was shown that when the number of scenarios is sufficiently large, the greedy and randomized methods produce comparable solutions, but in vastly different running times. The decreased running time of the randomized method allows for larger numbers of scenarios and, therefore, better solutions in practice.

One very important feature of the scenario approach is the distribution-free assumption of the underlying theorems. The absence of any distributional assumptions, such as finite moments and finite support, allows one to use the actual historical data (e.g., the Yahoo Finance data) as the scenarios, and then assess the properties of the obtained solutions by an application of the theorems without knowing the actual distribution that produced the data. In this paper, we pursued a different, indirect, approach by using data to first construct a probability distribution from which we sampled the scenarios. We did this because we wanted to test the performance of our method and to verify the correctness of the theoretical results. To improve the applicability of the method, future work includes devising scenario methods that exhibit lower scenario complexity. One promising line in this direction is represented by the introduction of a regularization term, possibly of $L1$ type, in the optimization procedure.

References

1. Nemirovski, A., Shapiro, A.: Convex approximations of chance constrained programs. SIAM J. Optim. **17**, 969–996 (2006)
2. Dentcheva, D., Prékopa, A., Ruszczynski, A.: Concavity and efficient points of discrete distributions in probabilistic programming. Math. Program. **89**, 55–77 (2000)
3. Prékopa, A.: Dual method for the solution of a one-stage stochastic programming problem with random rhs obeying a discrete probability distribution. Math. Methods Oper. Res. **34**, 441–461 (1990)
4. Tanner, M.W., Ntaimo, L.: IIS branch-and-cut for joint chance-constrained stochastic programs and application to optimal vaccine allocation. Eur. J. Oper. Res. **207**, 290–296 (2010)
5. Luedtke, J., Ahmed, S., Nemhauser, G.L.: An integer programming approach for linear programs with probabilistic constraints. Math. Program. **122**, 247–272 (2010)
6. Küçükyavuz, S.: On mixing sets arising in chance-constrained programming. Math. Program. **132**, 31–56 (2012)
7. Lejeune, M.: Pattern-based modeling and solution of probabilistically constrained optimization problems. SPEPS **2010**, 5 (2010). http://edoc.hu-berlin.de/browsing/speps/
8. Calafiore, G., Campi, M.C.: Uncertain convex programs: randomized solutions and confidence levels. Math. Program. **102**, 25–46 (2005)
9. Campi, M.C., Garatti, S.: The exact feasibility of randomized solutions of robust convex programs. SIAM J. Optim. **19**, 1211–1230 (2008)
10. Campi, M.C., Garatti, S.: A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality. J. Optim. Theory Appl. **148**, 257–280 (2010)
11. McNeil, A.J., Frey, R., Embrechts, P.: Quantitative Risk Management. Princeton University Press, Princeton (2005)
12. Pagnoncelli, B.K., Ahmed, S., Shapiro, A.: Computational study of a chance constrained portfolio selection problem. J. Optim. Theory Appl. **142**, 399–416 (2009)
13. Ben-Tal, A., Nemirovski, A.: Robust solutions of uncertain linear programs. Oper. Res. Lett. **25**, 1–14 (1999)
14. Shapiro, A., Dentcheva, D., Ruszczynski, A.: Lectures on Stochastic Programming: Modeling and Theory. SIAM, Philadelphia (2009)
15. Shapiro, A., Philpott, A.: A tutorial on stochastic programming. Manuscript. Available at www2.isye.gatech.edu/~ashapiro/publications.html (2007)
16. Wand, Y., Chen, Z., Zhang, K.: A chance-constrained portfolio selection problem under t-distribution. Asia-Pac. J. Oper. Res. **24**, 535–556 (2007)
17. Lejeune, M.: A var Black-Litterman model for the construction of absolute return fund-of-funds. Quant. Finance **11**, 1489–1501 (2011)
18. Nemirovski, A., Shapiro, A.: Scenario approximations of chance constraints. In: Calafiore, G., Dabbene, F. (eds.) Probabilistic and Randomized Methods for Design under Uncertainty, pp. 3–47. Springer, London (2006)
19. Luedtke, J., Ahmed, S.: A sample approximation approach for optimization with probabilistic constraints. SIAM J. Optim. **19**, 674–699 (2008)
20. Law, A., Kelton, W.D.: Simulation Modeling and Analysis. McGraw-Hill, New York (2000)