International Journal of Information Technology & Decision Making
Vol. 7, No. 4 (2008) 571–587

SIMULATION OPTIMIZATION: APPLICATIONS IN RISK MANAGEMENT

MARCO BETTER* and FRED GLOVER†
OptTek Systems, Inc., 1919 Seventh Street
Boulder, Colorado 80302, USA
**better@opttek.com*
†glover@opttek.com

GARY KOCHENBERGER
University of Colorado Denver, 1250 14th Street
Suite 215, Denver, Colorado 80202, USA
Gary.kochenberger@cudenver.edu

HAIBO WANG
Texas A&M International University
Laredo, TX 78041, USA
hwang@tamiu.edu

Simulation optimization is providing solutions to important practical problems previously beyond reach. This paper explores how new approaches are significantly expanding the power of simulation optimization for managing risk. Recent advances in simulation optimization technology are leading to new opportunities to solve problems more effectively. Specifically, in applications involving risk and uncertainty, simulation optimization surpasses the capabilities of other optimization methods not only in the quality of solutions but also in their interpretability and practicality. In this paper, we demonstrate the advantages of using a simulation optimization approach to tackle risky decisions, by showcasing the methodology on two popular applications from the areas of finance and business process design.

Keywords: Optimization; simulation; portfolio selection; risk management.

1. Introduction

Whenever uncertainty exists, there is risk. Uncertainty is present when there is a possibility that the outcome of a particular event will deviate from what is expected. In some cases, we can use past experience and other information to try to estimate the probability of occurrence of different events. This allows us to estimate a probability distribution for all possible events. *Risk* can be defined as the probability of occurrence of an event that would have a negative effect on a goal. On the other hand, the probability of occurrence of an event that would have a positive impact is considered an *opportunity* (see Ref. 1 for a detailed discussion of risks and opportunities). Therefore, the portion of the probability distribution that represents potentially harmful, or unwanted, outcomes is the focus of risk management.

Risk management is the process that involves identifying, selecting, and implementing measures that can be applied to mitigate risk in a particular situation.[1] The objective of risk management, in this context, is to find the set of actions (i.e. investments, policies, resource configurations, etc.) to reduce the level of risk to acceptable levels. What constitutes an acceptable level will depend on the situation, the decision makers' attitude toward risk, and the marginal rewards expected from taking on additional risk. In order to help risk managers achieve this objective, many techniques have been developed, both qualitative and quantitative. Among quantitative techniques, optimization has a natural appeal because it is based on objective mathematical formulations that usually output an optimal solution (i.e. set of decisions) for mitigating risk. However, traditional optimization approaches are prone to serious limitations.

In Sec. 2 of this paper, we briefly describe two prominent optimization techniques that are frequently used in risk management applications for their ability to handle uncertainty in the data; we then discuss the advantages and disadvantages of these methods. In Sec. 3, we discuss how simulation optimization can overcome the limitations of traditional optimization techniques, and we detail some innovative methods that make this a very useful, practical, and intuitive approach for risk management. Section 4 illustrates the advantages of simulation optimization on two practical examples. Finally, in Sec. 5 we summarize our results and conclusions.

2. Traditional Scenario-Based Optimization

Very few situations in the real world are completely devoid of risk. In fact, a person would be hard-pressed to recall a single decision in their life that was completely risk-free. In the world of deterministic optimization, we often choose to "ignore" uncertainty in order to come up with a unique and objective solution to a problem. However, in situations where uncertainty is at the core of the problem — as it is in risk management — a different strategy is required.

In the field of optimization, there are various approaches designed to cope with uncertainty.[2,3] In this context, the exact values of the parameters (e.g. the data) of the optimization problem are not known with absolute certainty, but may vary to a larger or lesser extent depending on the nature of the factors they represent. In other words, there may be many possible "realizations" of the parameters, each of which is a possible *scenario*.

Traditional scenario-based approaches to optimization, such as *scenario optimization* and *robust optimization*, are effective in finding a solution that is feasible for all the scenarios considered and minimizing the deviation of the overall solution from the optimal solution for each scenario. These approaches, however, only consider a very small subset of possible scenarios, and the size and complexity of models they can handle are very limited.

2.1. *Scenario optimization*

Dembo[4] offers an approach to solving stochastic programs based on a method for solving deterministic scenario subproblems and combining the optimal scenario solutions into a single feasible decision.

Imagine a situation in which we want to minimize the cost of producing a set J of finished goods. Each good j $(j = 1, \ldots, n)$ has a per-unit production cost c_j associated with it, as well as an associated utilization rate a_{ij} of resources for each finished good. In addition, the plant that produces the goods has a limited amount of each resource i $(i = 1, \ldots, m)$, denoted by b_i. We can formulate a deterministic mathematical program for a single scenario s (the scenario subproblem, or *SP*) as follows:

SP:

$$z^s = \text{minimize } \sum_{j=1}^{n} c_j^s x_j, \tag{1}$$

$$\text{Subject to } \sum_{j=1}^{n} a_{ij}^s x_j = b_i^s \quad \text{for } i = 1, \ldots, m, \tag{2}$$

$$x_j \geq 0 \quad \text{for } j = 1, \ldots, n, \tag{3}$$

where c^s, a^s, and b^s, respectively, represent the realization of the cost coefficient, the resource utilization, and the resource availability data under scenario s. Consider, for example, a company that manufactures a certain type of Maple door. Depending on the weather in the region where the wood for the doors is obtained, the costs of raw materials and transportation will vary. The company is also considering whether to expand production capacity at the facility where doors are manufactured, so that a total of six scenarios must be considered. The six possible scenarios and associated parameters for Maple doors are shown in Table 1. The first column corresponds to the particular scenario; Column 2 denotes whether the facility is at current or expanded capacity; Column 3 shows the probability of each capacity scenario; Column 4 denotes the weather (dry, normal, or wet) for each scenario; Column 5 provides the probability for each weather instance; Column 6 denotes the probability for each scenario; Column 7 shows the cost associated with each

Table 1. Possible scenarios for Maple doors.

Scen	Cap	$P(C)$	Wther	$P(W)$	P(Scen)	Cost c_j	Util a_{ij}	Avail b_j
1	Curr	50%	Dry	33%	1/6	L	H	L
2			Norm	33%	1/6	M	L	L
3			Wet	33%	1/6	H	L	L
4	Exp	50%	Dry	33%	1/6	L	H	H
5			Norm	33%	1/6	M	L	H
6			Wet	33%	1/6	H	L	H

scenario (L = low, M = medium, H = high); Column 8 denotes the utilization rate of the capacity (L = low, H = high); and Column 9 denotes the expected availability associated with each scenario.

The model *SP* needs to be solved once for each of the six scenarios. The scenario optimization approach can be summarized in two steps:

(1) Compute the optimal solution to each deterministic scenario subproblem *SP*.
(2) Solve a tracking model to find a single, feasible decision for all scenarios.

The key aspect of scenario optimization is the tracking model in step 2. For illustration purposes, we introduce a simple form of tracking model. Let p^s denote the estimated probability for the occurrence of scenario s. Then, a simple tracking model for our problem can be formulated as follows:

$$\text{Minimize} \quad \sum_s p^s \left(\sum_j c_j^s x_j - z^s \right)^2 + \sum_s p^s \left(\sum_{ij} a_{ij}^s x_j - b_i^s \right)^2, \tag{4}$$

$$\text{Subject to} \quad x_j \geq 0 \quad \text{for } j = 1, \ldots, n. \tag{5}$$

The purpose of this tracking model is to find a solution that is feasible under all the scenarios, and penalizes solutions that differ greatly from the optimal solution under each scenario. The two terms in the objective function are squared to ensure nonnegativity.

More sophisticated tracking models can be used for different purposes. In risk management, for instance, we may select a tracking model that is designed to penalize performance below a certain target level.

2.2. *Robust optimization*

Robust optimization may be used when the parameters of the optimization problem are known only within a finite set of values. The robust optimization framework gets its name because it seeks to identify a robust decision — i.e. a solution that performs well across many possible scenarios.

In order to measure the robustness of a given solution, different criteria may be used. Kouvelis and Yu identify three criteria: (1) absolute robustness; (2) robust deviation; and (3) relative robustness.[5] We illustrate the meaning and relevance of these criteria, by describing their robust optimization approach.

Consider an optimization problem where the objective is to minimize a certain performance measure such as cost. Let S denote the set of possible data scenarios over the planning horizon of interest. Also, let X denote the set of decision variables, and P the set of input parameters of our decision model. Correspondingly, let P^s identify the value of the parameters belonging to scenario s, and let F^s identify the set of feasible solutions to scenario s. The optimal solution to a specific scenario s

is then

$$z^s = f(X_s^*, P^s) = \min_{X \in F^s} f(X, P^s). \tag{6}$$

We assume here that f is convex. The first criterion, absolute robustness, also known as "worst-case optimization," seeks to find a solution that is feasible for all possible scenarios and optimal for the worst possible scenario. In other words, in a situation where the goal is to minimize the cost, the optimization procedure will seek the robust solution, z^R, that minimizes the cost of the maximum-cost scenario. We can formulate this as an objective function of the form

$$z^R = \min\{\max_{s \in S} f(X, P^s)\}. \tag{7}$$

Variations to this basic framework have been proposed (see, for example, Ref. 5) to capture the risk-averse nature of decision makers, by introducing higher moments of the distribution of z^s in the optimization model, and implementing weights as penalty factors for infeasibility of the robust solution with respect to certain scenarios.

The problem with both of these approaches, as with most traditional optimization techniques that attempt to deal with uncertainty, is their inability to handle a large number of possible scenarios. Thus, they often fail to consider events that, while unlikely, can be catastrophic. Recent approaches that use innovative simulation optimization techniques overcome these limitations by providing a practical, flexible framework for risk management and decision making under uncertainty.

3. Simulation Optimization

Simulation optimization can efficiently handle a much larger number of scenarios than traditional optimization approaches, as well as multiple sources and types of risk. Modern simulation optimization tools are designed to solve optimization problems of the form

$$\begin{aligned}
\text{Minimize } & F(x) \quad \text{(Objective function)},\\
\text{Subject to } & Ax \leq b \quad \text{(Constraints on input variables)},\\
& g_l \leq G(x) \leq g_u \quad \text{(Constraints on output measures)},\\
& l \leq x \leq u \quad \text{(Bounds)},
\end{aligned}$$

where the vector x of decision variables includes variables that range over continuous values and variables that only take on discrete values (both integer values and values with arbitrary step sizes).[6]

The objective function $F(x)$ is, typically, highly complex. Under the context of simulation optimization, $F(x)$ could represent, for example, the expected value of the probability distribution of the throughput at a factory; the fifth percentile of the distribution of the net present value (NPV) of a portfolio of investments; a measure of the likelihood that the cycle time of a process will be lower than a desired threshold value; etc. In general, $F(x)$ represents an output performance measure obtained from the simulation, and it is a mapping from a set of values x to a real value.

The constraints represented by inequality $Ax \leq b$ are usually linear (given that nonlinearity in the model is embedded within the simulation itself), and both the coefficient matrix A and the right-hand side values corresponding to vector b are known.

The constraints represented by inequalities of the form $g_l \leq G(x) \leq g_u$ impose simple upper and/or lower bound requirements on an output function $G(x)$ that can be linear or nonlinear. The values of the bounds g_l and g_u are known constants.

All decision variables x are bounded and some may be restricted to be discrete, as previously noted. Each evaluation of $F(x)$ and $G(x)$ requires an execution of a simulation of the system. By combining simulation and optimization, a powerful design tool results.

Simulation enables fast, inexpensive and nondisruptive examination and testing of a large number of scenarios prior to actually implementing a particular decision in the "real" environment. As such, it is quickly becoming a very popular tool in industry for conducting detailed "what-if" analysis. Since simulation approximates reality, it also permits the inclusion of various sources of uncertainty and variability into forecasts that impact performance. The need for optimization of simulation models arises when the analyst wants to find a set of model specifications (i.e. input parameters and/or structural assumptions) that leads to optimal performance. On the one hand, the range of parameter values and the number of parameter combinations are too large for analysts to enumerate and test all possible scenarios; hence, they need a way to *guide* the search for good solutions. On the other hand, without simulation, many real-world problems are too complex to be modeled by tractable mathematical formulations that are at the core of pure optimization methods like scenario optimization and robust optimization. This creates a conundrum; as shown above, pure optimization models alone are incapable of capturing all the complexities and dynamics of the system; hence, one must resort to simulation, which cannot easily find the best solutions. Simulation optimization resolves this conundrum by combining both methods.

Optimizers designed for simulation embody the principle of separating the method from the model. In such a context, the optimization problem is defined outside the complex system. Therefore, the evaluator (i.e. the simulation model) can change and evolve to incorporate additional elements of the complex system, while the optimization routines remain the same. Hence, there is a complete separation between the model that represents the system and the procedure that is used to solve optimization problems defined within this model.

The optimization procedure — usually based on metaheuristic search algorithms — uses the outputs from the system evaluator, which measures the merit of the inputs that were fed into the model. On the basis of both current and past

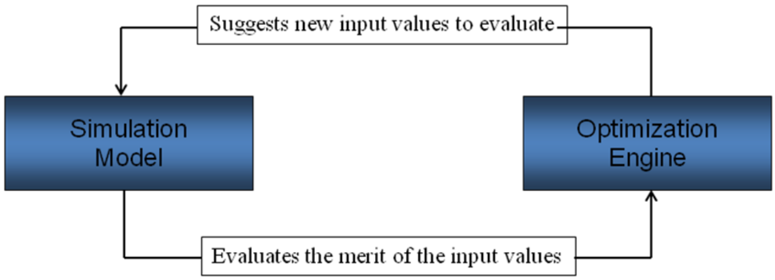

Fig. 1. Coordination between the optimization engine and the simulation.

evaluations, the method decides upon a new set of input values (Fig. 1 shows the coordination between the optimization engine and the simulation model).

Provided that a feasible solution exists, the optimization procedure ideally carries out a special search where the successively generated inputs produce varying evaluations, not all of them improving, but which over time provide a highly efficient trajectory to the globally best solutions. The process continues until an appropriate termination criterion is satisfied (usually based on the user's preference for the amount of time devoted to the search).

As stated before, the uncertainties and complexities modeled by the simulation are often such that the analyst has no idea about the shape of the response surface — i.e. the solution space. There exists no closed-form mathematical expression to represent the space, and there is no way to gauge whether the region being searched is smooth, discontinuous, etc. While this is enough to make most traditional optimization algorithms fail, metaheuristic optimization approaches, such as tabu search[7] and scatter search,[8] overcome this challenge by making use of adaptive memory techniques and population sampling methods that allow the search to be conducted on a wide area of the solution space, without getting stuck in local optima.

The metaheuristic-based simulation optimization framework is also very flexible in terms of the performance measures the decision-maker wishes to evaluate. In fact, the only limitation is not on the side of the optimization engine, but on the simulation model's ability to evaluate performance based on specified values for the decision variables. In order to provide in-depth insights into the use of simulation optimization in the context of risk management, we present some practical applications through the use of illustrative examples.

4. Illustrative Examples

4.1. *Selecting risk-efficient project portfolios*

Companies in the Petroleum and Energy (P&E) Industry use project portfolio optimization to manage investments in exploration and production, as well as power plant acquisitions.[9,10] Decision makers typically wish to maximize the return on invested capital, while controlling the exposure of their portfolio of projects to various risk factors that may ultimately result in financial losses.

In this example, we look at a P&E company that has 61 potential projects in its investment funnel. For each project, the pro-forma revenues for a horizon of 10–20 periods (depending on the project) are given as probability distributions. To carry it out, each project requires an initial investment and a certain number of business development, engineering and earth sciences personnel. The company has a budget limit for its investment opportunities, and a limited number of personnel of each skill category.

In addition, each project has a *probability of success* (POS) factor. This factor has a value between 0 and 1 and affects the simulation as follows: let us suppose that Project A has a POS = 0.6; therefore, during the simulation, we expect that we will be able to obtain the revenues from Project A in 60% of the trials, while in the remaining 40%, we will only incur the investment cost. The resulting probability distribution of results from simulating Project A would be similar to that shown in Fig. 2, where about 40% of the trials would have negative returns (i.e. equal to the investment cost), and the remaining 60% would have returns resembling the shape of its revenue distribution (i.e. equal to the simulated revenues minus the investment cost).

Projects may start in different time periods, but there is a restricted window of opportunity of up to three years for each project. The company must select a set of projects to invest in that will best further its corporate goals.

Probably, the best-known model for portfolio optimization is rooted in the work of Nobel laureate Harry Markowitz. Called the *mean-variance* model,[11] it is based on the assumption that the expected portfolio returns will be normally distributed. The model seeks to balance risk and return in a single objective function, as follows.

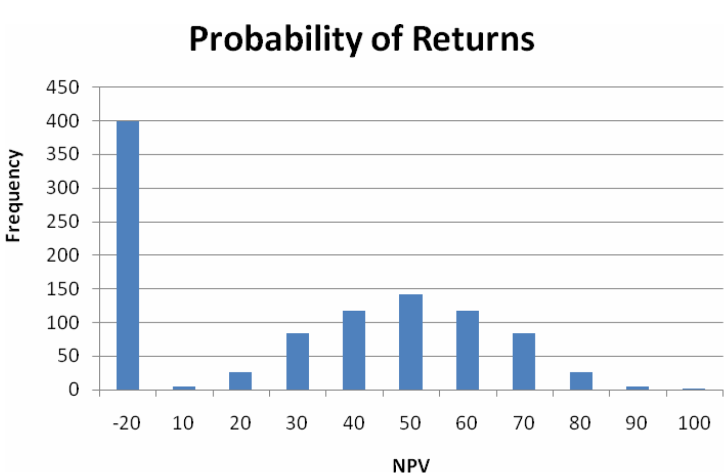

Fig. 2. Sample probability distribution of returns for a single simulated project.

Given a vector of portfolio returns r and a covariance matrix Q of returns, we can formulate the model as follows:

$$\text{Maximize} \quad r^{\mathrm{T}}w - kw^{\mathrm{T}}Qw, \tag{8}$$

$$\text{Subject to} \quad \sum_i c_i w_i = b, \tag{9}$$

$$w \in \{0, 1\}, \tag{10}$$

where k represents a coefficient of the firm's risk aversion, c_i represents the initial investment in project i, w_i is a binary variable representing the decision whether to invest in project i, and b is the available budget. We will use the mean-variance model as a base case for the purpose of comparing with other selected models of portfolio performance.

To facilitate our analysis, we make use of the OptFolio® software that combines simulation and optimization into a single system specifically designed for portfolio optimization.[12,13]

We examine three cases, including Value-at-Risk (VaR) minimization, to demonstrate the flexibility of this method to enable a variety of decision alternatives that significantly improve upon traditional mean-variance portfolio optimization, and illustrate the flexibility afforded by simulation optimization approaches in terms of controlling risk. The results also show the benefits of managing and efficiently allocating scarce resources like capital, personnel, and time. The weighted average cost of capital, or annual discount rate, used for all cases is 12%.

4.1.1. *Case 1: Mean-variance approach*

In this first case, we implement the mean-variance portfolio selection method of Markowitz described above. The decision is to determine participation levels (0 or 1) in each project with the objective of maximizing the expected NPV of the portfolio while keeping the standard deviation of the NPV below a specified threshold of \$140 M. We denote the expected value of the NPV by μ_{NPV}, and the standard deviation of the NPV by σ_{NPV}. This case can be formulated as follows:

$$\text{Maximize} \quad \mu_{\mathrm{NPV}} \quad \text{(objective function)}, \tag{11}$$

$$\text{Subject to} \quad \sigma_{\mathrm{NPV}} < \$140\mathrm{M} \quad \text{(requirement)}, \tag{12}$$

$$\sum_i c_i x_i \le b \quad \text{(budget constraint)}, \tag{13}$$

$$\sum_i p_{ij} x_i \le P_j \ \ \forall j \quad \text{(personnel constraints)}, \tag{14}$$

$$\text{All projects must start in year 1}, \tag{15}$$

$$x_i \in \{0, 1\} \ \ \forall i \quad \text{(binary decisions)}. \tag{16}$$

The optimal portfolio has the following performance metrics:

$$\mu_{\mathrm{NPV}} = \$394\mathrm{M}, \quad \sigma_{\mathrm{NPV}} = \$107\mathrm{M}, \quad P(5)_{\mathrm{NPV}} = \$176\mathrm{M},$$

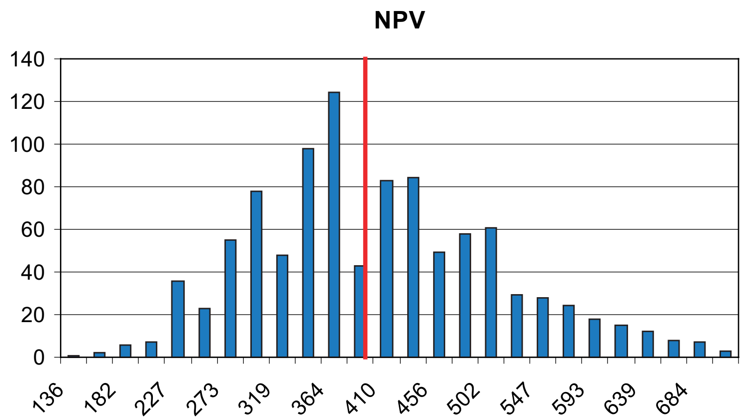

Fig. 2. Distribution of returns for mean-variance approach.

where $P(5)_{\mathrm{NPV}}$ denotes the fifth percentile of the resulting NPV probability distribution (i.e. the probability of the NPV being lower than the $P(5)$ value is 5%). The bound imposed on the standard deviation in Eq. (12) does not seem binding. However, due to the binary nature of the decision variables, no project additions are possible without violating the bound. Figure 2 shows a graph of the probability distribution of the NPV obtained from 1000 replications of this base model. The thin line represents the expected value.

4.1.2. *Case 2: Risk controlled by fifth percentile*

In the context of risk management, statistics such as variance or standard deviation of returns are not always easy to interpret, and there may be other measures that are more intuitive and useful. For example, it provides a clearer picture of the risk involved if we say "there is a 5% chance that the portfolio return will be below some value X" than to say that "the standard deviation is \$107M." The former analysis can be easily implemented in a simulation optimization approach by imposing a requirement on the fifth percentile of the resulting distribution of returns, as we describe here.

In Case 2, the decision is to determine participation levels (0, 1) in each project with the objective of maximizing the expected NPV of the portfolio, while keeping the fifth *percentile* of the NPV distribution above the value of the fifth percentile obtained in Case 1. In this way, we seek to "move" the distribution of returns further to the right, so as to reduce the likelihood of undesired outcomes.

This is achieved by imposing the requirement represented by Eq. (18) in the model below. In other words, we want to find the portfolio that produces the maximum average return, as long as no more than 5% of the trial observations fall below \$176M. In addition, we allow for delays in the start dates of projects, according to windows of opportunity defined for each project. In order to achieve this, in the simulation model we have created copies of each project that are shifted by one, two, or three periods into the future (according to the windows of opportunity defined for each project). Mutual exclusivity clauses among these stage copies of a project ensure that only one start date is selected. For example, to represent the fact that Project A can start at time $t = 0, 1$, or 2, we include the following mutual exclusivity clause as a constraint:

$$\text{Project A}_0 + \text{Project A}_1 + \text{Project A}_2 \leq 1.$$

The subscript following the project name corresponds to the allowed start dates for the project, and the constraint only allows at most one of these to be chosen:

$$\begin{aligned} \text{Maximize} \quad & \mu_{\text{NPV}}, && (17)\\ \text{Subject to} \quad & P(5)_{\text{NPV}} > \$176\text{M} \quad \text{(requirement)}, && (18)\\ & \sum_i c_i x_i \leq b \quad \text{(budget constraint)}, && (19)\\ & \sum_i p_{ij} x_i \leq P_j \ \ \forall j \quad \text{(personnel constraints)}, && (20)\\ & \sum_{m_i \in M} x_i \leq 1 \ \ \forall i \quad \text{(mutual exclusivity)}, && (21)\\ & x_i \in \{0, 1\} \ \ \forall I \quad \text{(binary decisions)}, && (22) \end{aligned}$$

where m_i denotes the set of mutually exclusive projects related to project i.

In this case we have replaced the standard deviation with the fifth percentile as a measure of risk containment. The resulting portfolio has the following attributes:

$$\mu_{\text{NPV}} = \$438\text{M}, \quad \sigma_{\text{NPV}} = \$140\text{M}, \quad P(5)_{\text{NPV}} = \$241\text{M}.$$

By using the fifth percentile instead of the standard deviation as a measure of risk, we are able to obtain an outcome that shifts the distribution of returns to the right, compared with Case 1, as shown in Fig. 3.

This case clearly outperforms Case 1. Although the distribution of returns exhibits a wider range (remember, we are not constraining the standard deviation here), not only we do obtain significantly better financial performance but we also achieve a higher personnel utilization rate and a more diverse portfolio.

4.1.3. *Case 3: Probability maximizing and VaR*

In Case 3, the decision is to determine participation levels $(0, 1)$ in each project with the objective of maximizing the probability of meeting or exceeding the mean NPV found in Case 1. This objective is expressed in Eq. (23) of the following model:

$$\text{Maximize} \quad P(\text{NPV} \geq \$394\text{M}), \tag{23}$$

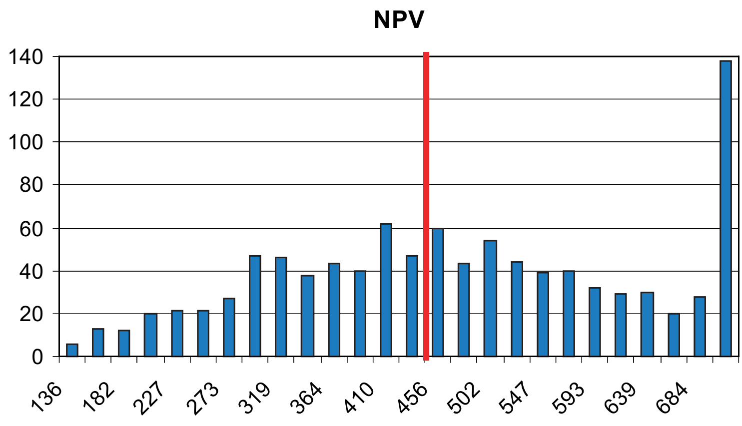

Fig. 3. Distribution of returns for Case 2.

$$\text{Subject to} \quad \sum_i c_i x_i \leq b \quad \text{(budget constraint)}, \tag{24}$$

$$\sum_i p_{ij} x_i \leq P_j \ \ \forall j \quad \text{(personnel constraints)}, \tag{25}$$

$$\sum_{m_i \in M} x_i \leq 1 \ \ \forall i \quad \text{(mutual exclusivity)}, \tag{26}$$

$$x_i \in \{0, 1\} \ \ \forall i \quad \text{(binary decisions)}. \tag{27}$$

This case focuses on maximizing the chance of achieving a goal and essentially combines performance and risk containment into one metric. The probability in (23) is not known *a priori*, hence, we must rely on the simulation to obtain it. The resulting optimal solution yields a portfolio that has the following attributes:

$$\mu_{\text{NPV}} = \$440\text{M}, \quad \sigma_{\text{NPV}} = \$167\text{M}, \quad P(5) = \$198\text{M}.$$

Although this portfolio has a performance similar to the one in Case 2, it has a 70% chance of achieving or exceeding the NPV goal (whereas Case 1 had only a 50% chance). As can be seen in the graph of Fig. 4, we have succeeded in shifting the probability distribution even further to the right, therefore increasing our chances of exceeding the returns obtained with the traditional Markowitz approach. In addition, in Cases 2 and 3, we need not make any assumption about the distribution of expected returns.

As a related corollary to the last case, we can conduct an interesting analysis that addresses VaR. In traditional (securities) portfolio management, VaR is defined as the worst expected loss under normal market conditions over a specific time interval and at a given confidence level. In other words, VaR measures how much the investor can lose, with probability $= \alpha$, over a certain time horizon.[14]

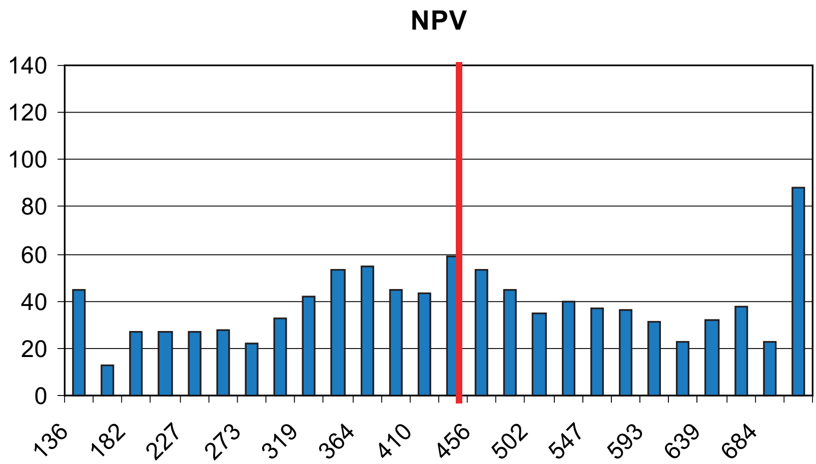

Fig. 4. Distribution of returns for Case 3.

In the case of project portfolios, we can define VaR as the probability that the NPV of the portfolio will fall below a specified value. Going back to our present case, the manager may want to limit the probability of incurring negative returns. In this example, we formulate the problem in a slightly different way: we still want to maximize the expected return, but we limit the probability that we incur a loss to $\alpha = 1\%$ by using the requirement shown in Eq. (29) as follows:

$$\text{Maximize} \quad \mu_{\text{NPV}}, \tag{28}$$

$$\text{Subject to} \quad P(\text{NPV} < 0) \leq 1\% \quad \text{(requirement)}, \tag{29}$$

$$\sum_i c_i x_i \leq b \quad \text{(budget constraint)}, \tag{30}$$

$$\sum_i p_{ij} x_i \leq P_j \ \ \forall j \quad \text{(personnel constraints)}, \tag{31}$$

$$\sum_{m_i \in M} x_i \leq 1 \ \ \forall i \quad \text{(mutual exclusivity)}, \tag{32}$$

$$x_i \in \{0, 1\} \ \ \forall i \quad \text{(binary decisions)}. \tag{33}$$

The portfolio performance under this scenario is

$$\mu_{\text{NPV}} = \$411\text{M}, \quad \sigma_{\text{NPV}} = \$159\text{M}, \quad P(5) = \$195\text{M}.$$

The results from the VaR model turn out to be slightly inferior to the case where the probability was maximized. This is not a surprise, since the focus of VaR is to limit the probability of downside risk, whereas before, the goal was to maximize the probability of obtaining a high expected return. However, this last analysis could prove valuable for a manager who wants to limit the VaR of the

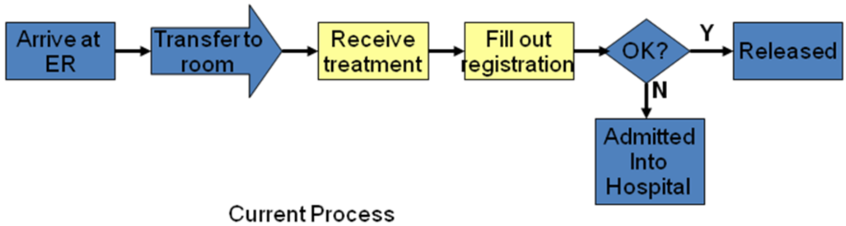

Fig. 5. Process flowchart for Level 1 patients.

selected portfolio. As shown here, for this particular set of projects, a very good portfolio — in financial terms — can still be selected with that objective in mind.

4.2. *Risk management in business process design*

Very common measures of process performance are the cycle time, a.k.a *turnaround time*, the throughput, and the operational cost. For our present example, we consider a process manager at a hospital emergency room (ER). Typically, emergency patients that arrive at the ER present different levels of criticality. In our example, we consider two levels: Level 1 patients are very critical and require immediate treatment; Level 2 patients are not as critical and must undergo an assessment by a triage nurse before being assigned to an ER room. Usually, a patient will have a choice in the care provider he or she prefers; hence, there is an inherent risk of lost business related to the quality of service provided. For instance, if patients must spend a very long time in the ER, there is a risk that patients will prefer to seek care at another facility on subsequent, follow-up visits.

Figure 5 shows a high-level flowchart of the process for Level 1 patients (the process for Level 2 patients differs only in that the "Fill out registration" activity is done before the "Transfer to room" activity, during the triage assessment).

The current operation consists of 7 nurses, 3 physicians, 4 patient care technicians (PCTs), 4 administrative clerks and 20 ER rooms. This current configuration has a total operating cost (i.e. wages, supplies, ER rooms' costs, etc.) of $52,600 per 100 h of operation, and the average time a Level 1 patient spends at the ER (i.e. cycle time) has been estimated at 1.98 h. Hospital management has just issued the new operational budget for the coming year and has allocated to the ER a maximum operational cost of $40,000 per 100 h of operation. The manager has three weeks to make changes to the process to ensure compliance with the budget, without deteriorating the current service levels.

The process manager has stated his goals as: *to minimize the average cycle time for Level 1 patients, while ensuring that the operational cost is below $40,000 per 100 h of operation.*

Since arrival times and service times in the process are stochastic, we use a simulation model to simulate the ER operation, and OptQuest to find the best configuration of resources in order to achieve the manager's goals. We obtain the

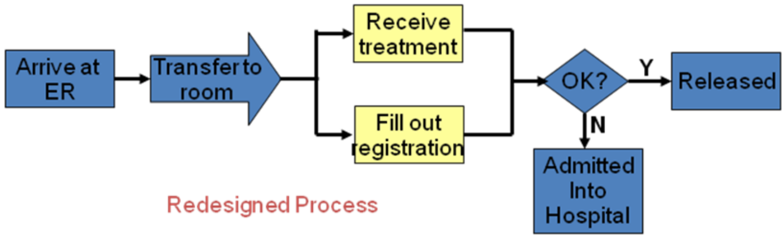

Fig. 6. Redesigned process for Level 1 patients.

following results:

- operational cost = \$36,200 (three nurses, three physicians, one PCT, two clerks, twelve ER rooms);
- average cycle time = 2.08 h.

The configuration of resources above (shown in parentheses) results in the lowest possible cycle time given the new operational budget for the ER. Obviously, in order to improve the service level at the current operational budget, it is necessary to re-design the process itself.

The new process proposed for Level 1 patients is depicted in Fig. 6. In the proposed process, the "Receive treatment" and "Fill out registration" activities are now done in parallel, instead of in sequence. The simulation of this new process, with the resource configuration found earlier, reduces the average cycle time from 2.08 to 1.98 h.

When designing business processes, however, it is very important to set goals that are not subject to the *Law of Averages* — that is, the goals should be set so that they accommodate a large percentage of the demand for the product or service that the process is intended to deliver. In the above example, it is possible that the probability distribution of cycle time turns out to be highly skewed to the right, so that even if the average cycle is under 2 h, there could be a very large number of patients who will spend a much longer time in the ER. It would be better to restate the manager's performance goals as follows: ensure that at least 95% of Level 1 patients will spend no more than 2 h in the ER *while ensuring that the operational cost is under \$40,000.* This gives a clear idea of the *service level* to which the process manager aspires.

If we re-optimize the configuration of resources with the new goal, we obtain the following results:

- operational cost = \$31,800 (four nurses, two physicians, two PCTs, two clerks, nine ER rooms);
- average cycle time = 1.94 h;
- 95th percentile of cycle time = 1.99 h.

Through the use of simulation optimization, we have obtained a process design that complies with the new budget requirements as well as with an improved service level. If implemented correctly, we can be sure that 95% of the critical patients in the ER will be either released or admitted into the hospital for further treatment in less than 2 h. Thus, the risk of having unsatisfied patients is minimized.

5. Results and Conclusions

Practically every real-world situation involves uncertainty and risk, creating a need for optimization methods that can handle uncertainty in model data and input parameters. We have briefly described two popular methods, *scenario optimization* and *robust optimization*, that seek to overcome limitations of classical optimization approaches for dealing with uncertainty and that undertake to find high-quality solutions that are feasible under as many scenarios as possible.

However, these methods are unable to handle problems involving moderately large numbers of decision variables and constraints, or involving significant degrees of uncertainty and complexity. In these cases, simulation optimization is becoming the method of choice. The combination of simulation and optimization affords all the flexibility of the simulation engine in terms of defining a variety of performance measures and risk profiles, as desired by the decision maker. In addition, as we demonstrate through two practical examples, modern optimization engines can enforce requirements on one or more outputs from the simulation, a feature that scenario-based methods cannot handle. This affords the user different alternatives for controlling risk, while ensuring that the performance of the system is optimized. The combination of simulation and optimization creates a tool for decision making that is fast (a simulation runs at a small fraction of real time, and the optimization guides the search for good solutions without the need to enumerate all possibilities), inexpensive, and nondisruptive (solutions can be evaluated without the need to stop the normal operation of the business, as opposed to pilot projects, which can also be quite expensive).

Finally, simulation optimization produces results that can be conveyed and grasped in an intuitive manner, providing the user with an especially useful and easy-to-use tool for identifying improved business decisions under risk and uncertainty.

References

1. D. Vose, *Risk Analysis: A Quantitative Guide* (John Wiley and Sons, Chichester, 2000).
2. M. Fukushima, How to deal with uncertainty in optimization — some recent attempts, *Int. J. Inform. Tech. Decis. Making* **5**(4) (2006) 623–637.
3. H. Eskandari and L. Rabelo, Handling uncertainty in the analytic hierarchy process: A stochastic approach, *Int. J. Inform. Tech. Decis. Making* **6**(1) (2007) 177–189.
4. R. Dembo, Scenario optimization, *Ann. Oper. Res.* **30** (1991) 63–80.

5. P. Kouvelis and G. Yu, *Robust Discrete Optimization and Its Applications* (Kluwer, Dordrecht, Netherlands, 1997), pp. 8–29.
6. J. P. Kelly, Simulation optimization is evolving, *INFORMS J. Comput.* **14**(3) (2002) 223–225.
7. F. Glover and M. Laguna, *Tabu Search* (Kluwer, Norwell, MA, 1997).
8. F. Glover, M. Laguna and R. Martí, Fundamentals of scatter search and path relinking, *Control Cybernet.* **29**(3) (2000) 653–684.
9. W. J. Haskett, Optimal appraisal well location through efficient uncertainty reduction and value of information techniques, in *Proc. Society of Petroleum Engineers Annual Technical Conference and Exhibition* (Denver, CO, 2003).
10. W. J. Haskett, M. Better and J. April, Practical optimization: Dealing with the realities of decision management, in *Proc. Society of Petroleum Engineers Annual Technical Conference and Exhibition* (Houston, TX, 2004).
11. H. Markowitz, Portfolio selection, *J. Finance* **7**(1) (1952) 77–91.
12. J. April, F. Glover and J. P. Kelly, Portfolio optimization for capital investment projects, in *Proc. 2002 Winter Simulation Conference*, eds. S. Chick, T. Sanchez, D. Ferrin and D. Morrice (2002), pp. 1546–1554.
13. J. April, F. Glover and J. P. Kelly, Optfolio — A simulation optimization system for project portfolio planning, in *Proc. 2003 Winter Simulation Conference*, eds. S. Chick, T. Sanchez, D. Ferrin and D. Morrice (2003), pp. 301–309.
14. S. Benninga and Z. Wiener, Value-at-risk (VaR), *Math. Educ. Res.* **7**(4) (1998) 1–7.