

Simulation Optimization:
Improving decisions under uncertainty

Dr. Marco Better

better@opttek.com

Dr. Fred Glover

glover@opttek.com

OptTek Systems, Inc.

Dr. Gary Kochenberger

gary.kochenberger@ucdenver.edu

University of Colorado Denver

INTRODUCTION

Analytics has been defined as “the scientific process of transforming data into insight for **making better decisions.**” More and more organizations are using analytics to make better decisions and reduce risk. Analytics includes well-established methods such as mathematical optimization, simulation, probability theory and statistics, as well as newer techniques that take elements from traditional methods and modify and/or combine them into robust frameworks in order to develop more powerful solution methods for many settings where traditional methods fall short. A prime example of the latter is the framework known as Simulation Optimization. As its name implies, this method combines simulation and optimization in order to tackle complex situations where risk and uncertainty do not behave according to certain simplifying assumptions.

Taken separately, each method is critical, but limited in scope. On the one hand, optimization by itself provides an excellent method to select the best element in terms of some system performance criteria, from some set of available alternatives, in the absence of uncertainty. On the other hand, simulation is a tool that allows us to build a representation of a complex system in order to better understand the uncertainty in the system’s performance.

By putting these two methods together, we can develop a powerful framework that takes advantage of each method’s strengths, so that we have at our disposal a technique that allows us to select the best element from a set of alternatives and simultaneously take account of the uncertainty in the system.

In this chapter, we will begin with an example to illustrate each technique separately, and highlight the benefits of simulation optimization approaches in the presence of uncertainty. We will then explore the use of simulation optimization in real-world applications in risk-management. Finally, we will summarize our discussion in a set of concluding remarks.

AN ILLUSTRATIVE EXAMPLE

“Portfolio investment theory” is concerned with finding the portfolio of investments that maximizes expected returns while minimizing risk. Since investors are risk-averse, they prefer portfolios with high expected returns and low risk. [1]

In 1952, Nobel laureate Dr. Harry Markowitz laid down the foundation for modern investment theory. Markowitz focused his attention on *mean-variance efficient* portfolios. [2] A portfolio is *mean-variance efficient* if it has the highest expected (mean) return for a given variance, or, similarly, if it has the smallest variance for a given expected return. Markowitz developed this theory for portfolios of securities, such as stocks and bonds, for which returns are usually normally distributed. If portfolio returns are normally distributed, then its risk can be measured by the variance of its returns. If this is the case, then the *optimal* set of portfolios can be found by traditional optimization methods, as we show in the following section.

Optimization of securities portfolios

What constitutes the best portfolio of securities? In 1952, Markowitz attempted to answer this question with his famous paper, "Portfolio Selection." In this paper, Markowitz formulated the portfolio selection problem as an optimization problem.

For simplicity, let's first consider a market that contains only three assets (i.e., stocks), A_1 , A_2 , and A_3 . Let's assume that we have a limited budget to invest in these assets, and that we want to invest our budget in its entirety among these assets. Therefore, we will denote x_1 , x_2 , and x_3 as the proportion of our budget that we will invest in A_1 , A_2 , and A_3 , respectively. Since we will invest our entire budget, it follows that:

$$x_1 + x_2 + x_3 = 1$$

Now, let μ_1 , μ_2 , and μ_3 denote the expected value (i.e., the mean) of the return of A_1 , A_2 , and A_3 , respectively; and let $\sigma_{(1)}^2$, $\sigma_{(2)}^2$, and $\sigma_{(3)}^2$ denote the variance of the probability distribution of the returns of A_1 , A_2 , and A_3 , respectively. Thus, we will use $\sigma_{(1,2)}^2$ to represent the covariance between the probability distributions of the returns of A_1 , A_2 , and so forth.

The mean returns, return variances and covariances can be estimated from historic stock price data.

The mean return of the portfolio, μ_p , will be equal to the weighted average expected return, where the weights correspond to the proportion of the budget invested in each asset; thus,

$$\mu_p = \mu_1 x_1 + \mu_2 x_2 + \mu_3 x_3$$

The variance of the probability distribution of the portfolio returns, $\sigma_{(p)}^2$, is calculated as:

$$\sigma_{(p)}^2 = \sigma_{(1,2)}^2 x_1 x_2 + \sigma_{(1,3)}^2 x_1 x_3 + \sigma_{(2,3)}^2 x_2 x_3$$

Therefore, we can formulate the problem of finding the best portfolio as:

$$\text{Maximize} \quad (\mu_p - \sigma_{(p)}^2) \quad (1)$$

$$\text{Subject to:} \quad x_1 + x_2 + x_3 = 1 \quad (2)$$

$$x_1, x_2, x_3 \geq 0 \quad (3)$$

Equation (1) is known as the objective function, because it represents our main goal, or objective. In this case, our objective is to find the portfolio with the maximum expected return and minimum risk.

Equation (2) makes sure the entire budget is invested among the assets; and Equation (3) enforces a positivity constraint on our investments (no short positions are allowed).

We could have used the alternate objective function:

$$\text{Minimize } (\sigma_{(p)}^2 - \mu_p) \quad (1')$$

In this case, our objective is to find the portfolio with minimum risk and maximum return. It is apparent that Equations (1) and (1') are equivalent.

These formulations can be solved using a specific type of mathematical programming technique called “quadratic programming.” By using quadratic programming, we can find the optimal solution(s) to the above portfolio selection problem. Unfortunately, there are two complicating factors here:

1. There may be many – in certain cases, even an infinite number of –optimal solutions to the above formulation. In fact, Markowitz calls the set of solutions the “efficient set,” and they all lie along a curve called the “efficient frontier.” It is up to the investor to “pick” the solution on the frontier that provides maximum return for the level of risk she is willing to accept.
2. The above formulation is valid only if one very strict assumption holds: the returns of each of the assets, and, hence, the portfolio returns must follow a normal probability distribution. If this is not the case, then the Markowitz model breaks down. Under these strict normality conditions the variance is symmetric, so that the probability that the actual portfolio return will be above its estimated mean is the same as the probability that it will be below. However, using the variance as the only risk measure does not protect us from the probability that the portfolio return will be very small, or even negative; in other words, we have no idea of the extent to which our investment is “at risk.”

In most practical situations, especially when there is a large number of underlying assets to be considered, these complicating factors make any solution found by this method unreliable. On one hand, it is difficult to quantify the exact level of risk a particular investor is willing to accept. On the other hand, portfolio returns often do not follow a normal distribution, or the distribution is severely skewed when investment costs and capital gains tax implications come into play. Therefore, although we have a very elegant model like Markowitz’s, we cannot use it.

We need a method that can provide us with more complete information about the quality of the portfolio in terms of risk and return.

Simulation

Monte Carlo simulation is a method used by financial companies to simulate and understand risks related to various investments. [3] The main advantage of this method is that the normality assumption is no longer a requirement; in fact, the power of the method is that we can use statistical techniques to analyze an asset’s historical data, and forecast its behavior into the future

by simulating the probable outcomes. This provides freedom from strict assumptions about the probability distribution of the assets.

To illustrate this better, we have taken a sample of month-end closing stock prices for six very well-known high-tech corporations: Sun Microsystems, Oracle, Microsoft, Intel Corp., Cisco Systems, and Yahoo, for the period between March 31, 1999 and April 30, 2001.

Let's take for instance, the end-of-month closing stock price for Intel Corporation (INTEL). During the (roughly) two-year period INTEL's end-of-month stock price averaged \$75.50, with an average monthly change of -2.74 percentage points, and an overall drop of 74% from \$118.88 to \$30.91. During that period, the month-end stock price increased 11 times and decreased 14 times. The biggest month-to-month change for INTEL during that period was a 48.5% drop, from a \$118.88 closing price on March 31, 1999 to April 30, 1999; conversely, the biggest increase of 23.1%, from a closing price of \$30.03 to \$37.00 was registered between December 30, 2000 and January 30, 2001.

Using this information and some additional statistics, we can find a probability distribution that these data fits quite well. Several commercial software products exist that can perform "goodness-of-fit" testing for data such as these. For our example, we used Crystal Ball's "Batch Fit" feature, which automatically finds the best fit among a set of more than a dozen well-known probability distributions (see <http://www.oracle.com/us/products/applications/crystalball> for more information about the Batch Fit and other features).

According to Crystal Ball, the historic end-of-month stock price best fits a Lognormal probability distribution with Mean=\$76.03 and Standard Deviation=\$39.22. If we plot this distribution, we get a graph like the one shown in Figure 1.

From Figure 1 we can deduce that there is an 80% likelihood that month-end stock prices for INTEL will fall between roughly \$33.00 and \$127.00, with only a 20% chance of being higher or lower.




Figure 1: Crystal Ball screen shot of lognormal probability distribution

Similarly, we can find the best-fitting distributions for each of the stocks that we are interested in including in our portfolio. These results are summarized in Table 1.

Table 1: summary statistics for selected stock prices

Stock	Mean Stock Price	Standard Deviation	P(10)	P(90)	Selected Distribution	Average Monthly Return
SUN	\$77.96	\$33.84	\$30.99	\$117.57	Min Extreme	-3.69%
ORACLE	49.93	26.85	18.44	89.90	Beta	2.93%
MICROSOFT	78.33	17.01	56.59	100.62	Beta	0.39%
INTEL	76.03	35.79	33.82	127.04	Lognormal	-2.74%
CISCO	68.80	30.52	30.81	108.86	Lognormal	-5.18%
YAHOO	130.30	93.48	21.68	238.91	Logistic	-3.74%

Based on information contained in the third column and the last column of Table 1 which contain the standard deviation and the average monthly return for each stock, if we were to use the pure optimization approach discussed previously, the optimal portfolio would result in an average monthly return of 2.59% over the next two years, with a standard deviation of returns of 0.19%, by taking the following portfolio positions:

- Invest about 25.7% of our budget in ORACLE stock;
- Invest about 17.7% of the budget in MICROSOFT stock;
- Invest about 47.5% of the budget in INTEL stock; and,
- Invest the remaining 9.1% of the budget in CISCO stock.

However, if we simulate this portfolio using Monte Carlo simulation and the best-fit distribution information above, we obtain the following results:

- Average monthly portfolio return = 0.99%

- Standard deviation of returns = 0.93%

We can see that the simulation of the portfolio, which takes into account the uncertainty in the stock prices (i.e., the variability of stock prices due to a number of known and unknown factors), results in a much more conservative performance. In fact, according to the model there is an almost 10% chance that the return on this portfolio will be negative.

So, can we pick a portfolio that is better? Well, what if we just invested an equal amount in each stock? We can simulate this portfolio with a $1/6 = 16.66\%$ investment of the budget in each stock. The expected value of the monthly return based on point estimates is 0.63%; however, the simulation results in an average monthly return of 0.97%, with a standard deviation of 0.49%. Although the expected return is slightly lower than for the previous portfolio, the standard deviation has been greatly reduced, so this can be considered a “safer” portfolio. In fact this portfolio has only a 1.55% chance of resulting in a negative return.

The question is: How can we find the best portfolio given the uncertainty in stock prices? We quickly realize that it is extremely unlikely that we will find it by “trying out” different portfolio alternatives by hand, since there are so many possible combinations of assets and budget allocations. So we need something more powerful to help us.

The solution to this challenge: combine the advantages of optimization and simulation into a single framework. On one hand, optimization can help us search for the best portfolio; on the other, simulation can ensure that we are not ignoring the uncertainty in stock prices.

A Simulation Optimization solution approach

As we learned in the previous section, simulation provides a way to evaluate the impact of changes to parameters and decisions in a model environment through the creation of “what-if” scenarios. Simulation also enables examination and testing of decisions prior to actually implementing them, thus saving the decision maker time and money. But perhaps most important, simulation enables the examination and evaluation of uncertainty and risks associated with a set of decisions, so that such risks can be understood and mitigated.

Although simulation provides all of these advantages, in most real-world situations – even one as simple as our portfolio example – the set of possible decisions is too large to enumerate, let alone to be searched through efficiently to find the best possible solution. We need a way to guide our search.

The merging of simulation modeling with optimization technology has provided the advance making it possible to solve this problem. [4]

Once a simulation model has been developed to represent a system or process, we want to find a configuration that is *best* according to some performance criteria. In our portfolio example, this corresponds to finding investment levels in the set of six candidate stocks that will result in

maximum expected returns at some minimum risk level. When a system is complex and the configuration depends on a number of strategic choices, the trial-and-error approach results in very limited success. In these cases, we use an optimization tool to guide the search for the best configuration.

In our approach, we view the simulation model as a “black box”, meaning that we are only interested in obtaining an evaluation of performance from the simulation.

Figure 2 shows this black-box approach to simulation optimization. In this approach, the optimization procedure, called the *optimizer*, first chooses a set of values for the input parameters, also called the *decision variables*; next, the simulation model is run with these parameter values, and the optimizer obtains an *evaluation* of the performance of the system. The optimizer then uses this response from the simulation model to select the next set of values for the input parameters. This loop continues until certain stopping rules are met.

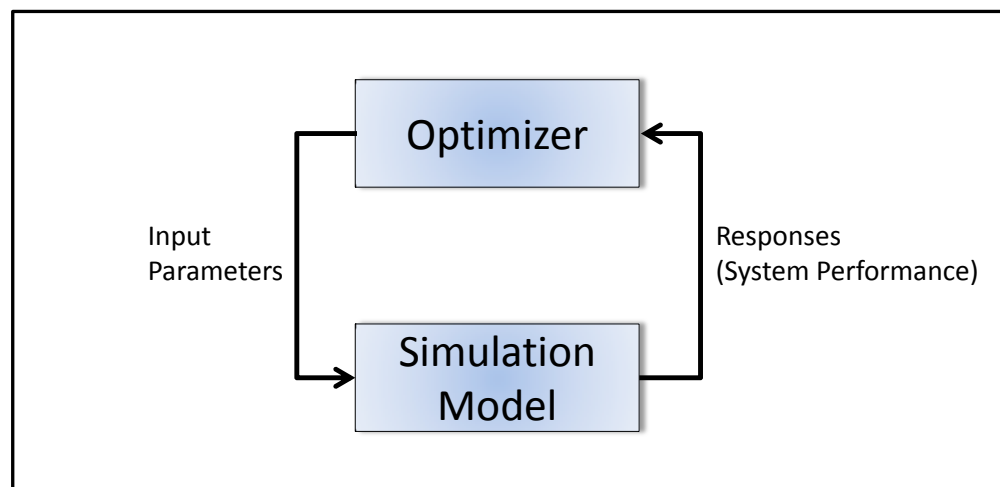


Figure 2: black box approach to simulation optimization

The key to a good simulation optimization approach is the search algorithm embedded in the optimizer, which makes decisions about the selection of new input parameter values at each iteration in the process.

Going back to our portfolio selection example, we now view our simulation model as a black box that provides a response about the expected return and standard deviation of returns of a selected portfolio. We wrap around it an optimizer that will guide us in selecting the allocation of our budget to each stock in order to obtain the maximum expected return at some minimum level of risk.

In order to do this, we need to tell the optimizer exactly what we want to achieve. Therefore, we first state our primary objective: *maximize the expected portfolio return*. If we run the optimization with this objective alone, we obtain the following portfolio recommendation:

- Invest the entire budget in ORACLE stock.

This is expected to produce the following results:

- Average monthly portfolio return = 1.06%
- Standard deviation of returns = 4.08%

In this case, since we did not tell the optimizer anything about the risk we are willing to accept, the optimizer naturally chose to invest the entire budget in the stock with the highest expected return, regardless of its volatility. We want to do better than this. So we amend our objective to include the additional goal: keep risk under a certain acceptable threshold.

We now express our desires as follows:

Maximize the expected portfolio return, but make sure the standard deviation of returns is no higher than 0.2%.

If we optimize this model, the best solution found – at iteration 861 – recommends the following:

- 11% of budget should be invested in SUN MICROSYSTEMS stock;
- 46% of budget should be invested in MICROSOFT stock;
- 28% of budget should be invested in INTEL stock;
- 15% of budget should be invested in CISCO stock.

This yields the following results:

- Average monthly portfolio return = 0.96%
- Standard deviation of returns = 0.19%

It is notable that the optimal portfolio in this case includes NO investment in ORACLE stock in spite of it having the highest historical return. This is probably because ORACLE has the highest ratio of standard deviation to average return, making it the most volatile of the stocks, and thus increasing the riskiness of the portfolio. In fact, according to our amended model, the optimal portfolio described above has only a 0.04% chance of resulting in a negative return, which is a vast improvement.

SIMULATION OPTIMIZATION applications in other real-world settings

The advantages of simulation optimization can not only be realized in financial modeling settings. A wide array of fields exists where systems exhibit high complexity and outcomes are sensitive to uncertainty. Examples are business processes, national defense systems, workforce planning, and so forth. When changes are proposed to a system in order to improve performance, the projected improvements can be simulated and optimized artificially. The sensitivity of performance objectives to proposed changes can be examined and quantified, reducing the risk of actual implementation, and increasing the confidence in the selected decision strategy.

In business process management, changes may entail adding, deleting, and modifying processes, process times, resource requirements, schedules, work rates, skill levels, and budgets, making this a very fertile area for such approaches. Performance objectives may include throughput, costs, inventories, cycle times, resource and capital utilization, start-up and set-up times, cash flow and waste. In the context of business process management and improvement, simulation can be thought of as a way to understand and communicate the uncertainty related to making the changes while optimization provides the way to manage that uncertainty.

We now focus on two examples to showcase the application and advantages of a simulation optimization approach to business process management.

1 SELECTING THE BEST CONFIGURATION IN A HOSPITAL EMERGENCY ROOM

The following example is based on a model of a real emergency room (ER) process provided by CACI, and simulated on SIMPROCESS. Consider the operation of an emergency room in a hospital. Figure 1 shows a high-level view of the overall process. The process begins when a patient arrives through the doors of the ER, and ends when a patient is either released from the ER or admitted into the hospital for further treatment. Upon arrival, patients sign in, are assessed in terms of their condition, and are transferred to an ER room. Depending on their condition, patients must then go through the registration process and through the treatment process, before being released or admitted into the hospital.

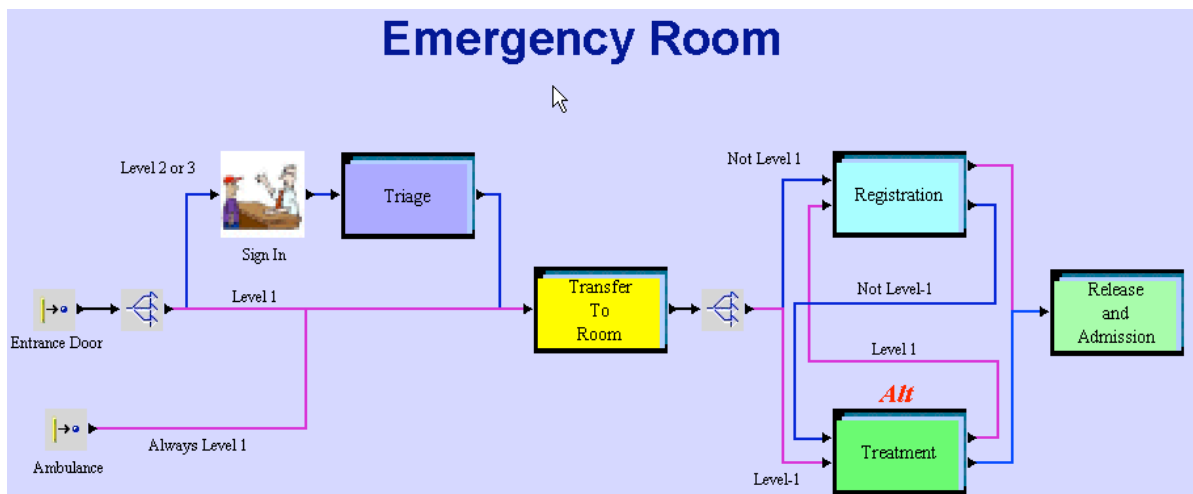


Figure 3: High-Level Process View

Patients arrive either on their own or in an ambulance, according to some arrival process. Arriving patients are classified into different levels, based on their condition, with Level 1 patients being more critical than Level 2 and Level 3.

Level 1 patients are taken to an ER room immediately upon arrival. Once in the room, they undergo their treatment. Finally, they complete the registration documentation process before being either released or admitted into the hospital for further treatment.

Level 2 and Level 3 patients must first sign in with an *Administrative Clerk*. After signing in, their condition is assessed by a *Triage Nurse*, and then they are taken to an ER room. Once in the room, Level 2 and 3 patients, must first complete their registration documents, then go on to receive their treatment, and, finally, they are either released or admitted into the hospital for further treatment.

The ER treatment process consists of the following activities:

1. A secondary assessment performed by a nurse and a physician.
2. Laboratory tests, if necessary, performed by a patient care technician (PCT).
3. The treatment itself, performed by a nurse and a physician.

The registration process consists of the following activities:

1. A data collection activity performed by an Administrative Clerk.
2. An additional data collection activity performed by an Administrative Clerk, in case the patient has Worker's Compensation Insurance.
3. A printing of the patient's medical chart for future reference, also performed by an Administrative Clerk.

Finally, 90% of all patients are released from the ER, while the remaining 10% are admitted into the hospital for further treatment. The final release/hospital admission process consists of the following activities:

1. In case of release, either a nurse or a PCT fills out the release papers (based on availability).
2. In case of admission into the hospital, an Administrative Clerk fills out the patient's admission papers. The patient must then wait for a hospital bed to become available. The time until a bed is available is handled by an empirical probability distribution. Finally, the patient is transferred to the hospital bed.

The ER has the following resources:

- Nurses
- Physicians
- PCTs

- Administrative Clerks
- ER Rooms

In addition, the ER has one Triage Nurse and one Charge Nurse at all times.

Due to cost and layout considerations, hospital administrators have determined that the staffing level must not exceed 7 nurses, 3 physicians, 4 PCTs and 4 Administrative Clerks. Furthermore, the ER has 20 rooms available; however, using fewer rooms would be beneficial, since the additional space could be used more profitably by other departments in the hospital. The hospital wants to find the configuration of the above resources that minimizes the total asset cost.

The asset cost includes the staff's hourly wages and the fixed cost of each ER room used. We must also make sure that, on average, Level 1 patients do not spend more than 2.4 hours in the ER. This can be formulated as an optimization problem, as follows:

Minimize the expected Total Asset Cost

Subject to the following constraints:

- Average Level 1 Cycle Time is less than or equal to 2.4 hours
- # Nurses are greater than or equal to 1 and less than or equal to 7
- # Physicians are greater than or equal to 1 and less than or equal to 3
- # PCT's are greater than or equal to 1 and less than or equal to 4
- # Admin. Clerks are greater than or equal to 1 and less than or equal to 4
- # ER Rooms are greater than or equal to 1 and less than or equal to 20

This is a relatively simple problem in terms of size: 6 variables and 6 constraints. However, if we were to rely solely on simulation to solve this problem, even after the hospital administrators have narrowed down our choices to the above limits, we would have to perform $7 \times 3 \times 4 \times 4 \times 20 = 6,720$ experiments. If we want a sample size of, say, at least 30 runs per trial solution in order to obtain the desired level of precision, then each experiment would take about 2 minutes¹. This means that a complete enumeration of all possible solutions would take approximately 13,400 minutes, or about 28 working days. This is obviously too long a duration for finding a solution.

In order to solve this problem in a reasonable amount of time, we used the OptQuest[®] optimization technology integrated with SIMPROCESS (see www.OptTek.com for more information about the OptQuest optimizer). As a base case we decided to use the upper resource limits provided by hospital administrators, to get a reasonably good initial solution. This

¹ We timed one experiment with 30 runs on a Dell Dimension 8100, with an Intel Pentium 4 processor @ 1700 MHz.

configuration yielded an Expected Total Asset Cost of \$ 36,840, and a Level 1 patient cycle time of 1.91 hours.

Once we set up the problem in OptQuest, we ran it for 100 iterations (experiments), and 5 runs per iteration (each run simulates 5 days of the ER operation). Given these parameters, the best solution, found at iteration 21 was:

Nurses	Physicians	PCT's	Admin Clerks	ER Rooms
4	2	3	3	12

The Expected Total Asset Cost for this configuration came out to \$ 25,250 (a 31% improvement over the base case), and the average Level 1 patient cycle time was 2.17 hours. However, looking at the probability distribution for cycle time, we see that there is still a 45% chance that the Level 1 patient cycle time will be greater than 2.40 hours.

After obtaining this solution, we redesigned some features of the current model to improve the cycle time of Level 1 patients even further. In the proposed model, we assume that Level 1 patients can go through the treatment process and the registration process in parallel. That is, we assume that while the patient is undergoing treatment, the registration process is being done by a surrogate or whoever is accompanying the patient. If the patient's condition is very critical, than someone else can provide the registration data; however, if the patient's condition allows it, then the patient can provide the registration data during treatment. Figure 2 shows the model with this change. By implementing this change in the optimized model, we now obtain an average Level 1 patient cycle time of 1.98 (a 12% improvement), with only a 30% chance that the Level 1 patient cycle time will be greater than 2.40 hours.

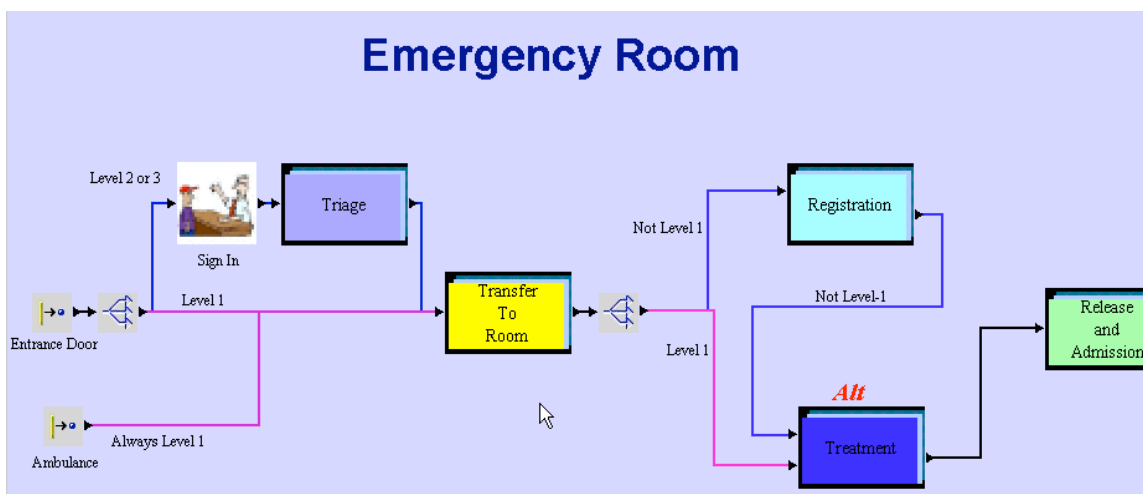


Figure 4: Proposed Process

Upon re-optimizing this new model, given the change that we implemented, we obtain a new optimal solution in 28 iterations, as follows:

Nurses	Physicians	PCT's	Admin Clerks	ER Rooms
4	2	2	2	9

This configuration yields an Expected Total Asset Cost of \$ 24,574, and an average Level 1 patient cycle time of 1.94 hours, with less than a 5% probability that the cycle time will exceed 2.40 hours. By using optimization, we were able to find a very high quality solution in less than 30 minutes, which provides a very acceptable level of risk in terms of service quality (i.e., cycle time). In addition, we were able to make changes to improve the model and re-optimize to find a better configuration. It is highly unlikely that this solution would be found relying solely on simulation.

2 SELECTING THE BEST STAFFING LEVEL FOR A PERSONAL CLAIMS PROCESS AT AN INSURANCE COMPANY

The following example is based on a model provided by SIMUL8 Corporation. We used the *SIMUL8* simulation tool for the simulation, and *OptQuest for SIMUL8* for the optimization.

A personal claims department in an insurance company handles claims made by their clients. Claims arrive according to a Poisson process, with a mean inter-arrival time of 5 minutes. Figure 3 is a process map depicting the personal claims process in terms of swim lanes.

The first lane corresponds to work done by a Claims Handler (CH) located at the client's local service center. Upon arrival of a claim, an Assessor determines if the client has a valid policy. If not (5% of all cases), then the case is terminated; otherwise (95% of all cases), the Assessor enters the appropriate information in the system.

In the second lane, an Assessor located at the Service Center (ASC) receives the information from the Claims Handler. The Assessor first determines if the claim is covered by the client's policy. If not (5% of all cases), the case is terminated; otherwise (95% of all cases), the Assessor approves the preliminary estimate of the damage. If the damage exceeds \$2,000 (35% of all cases), the claim is sent to an Assessor at Headquarters for approval; otherwise (65% of all cases), it is sent directly to a Senior Assessor.

Lane 3 corresponds to the Assessor at Headquarters (AHQ). The Assessor first authorizes the on-site investigation of the accident. If the investigation determines that the incident is not covered by the client's policy (2% of all cases), then the case is terminated; otherwise (98% of all cases), a final price is determined and the case is approved.

In lane 4, the Senior Assessor (SA) receives the claim, checks it for completeness, and provides the final approval. Once the claim is approved it is sent to Document Control.

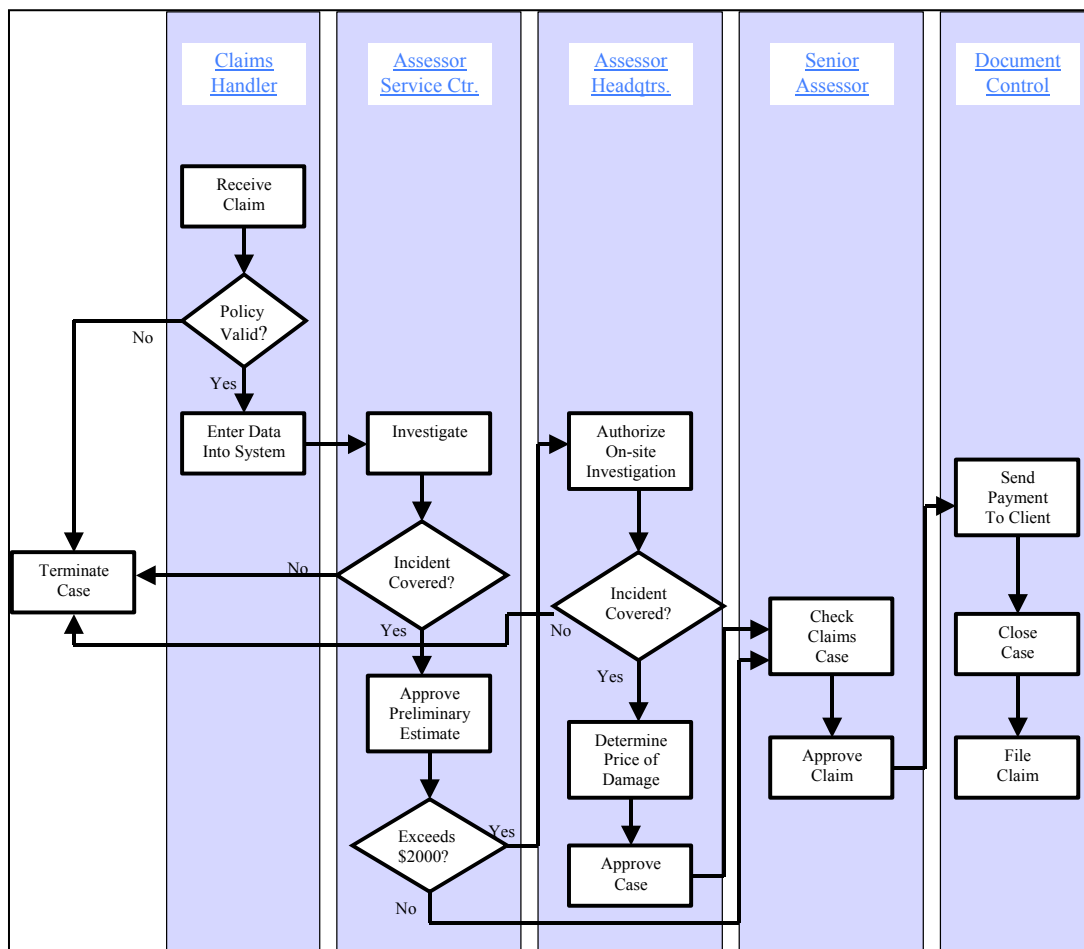


Figure 5: Map of Personal Claims Process

Document Control (DC), in lane 5, is in charge of processing the payment to the client, closing the case and, finally, filing the claim.

The objective here is to find staffing levels for each of the five resource types, in order to minimize headcount, while keeping average throughput above 1,500 claims during 4 weeks. Each resource type has a maximum limit of 20 people, and the overall headcount in the process cannot exceed 90. The formulation of the optimization problem is as follows:

Minimize the Headcount

Subject to the following constraints:

- Average Throughput is equal to or greater than 1500

- # Claims Handlers are greater than or equal to 1 and less than or equal to 20
- # Service Center Assessors are greater than or equal to 1 and less than or equal to 20
- # Headquarter Assessors are greater than or equal to 1 and less than or equal to 20
- # Senior Assessors are greater than or equal to 1 and less than or equal to 20
- # Document Controllers are greater than or equal to 1 and less than or equal to 20
- The overall headcount cannot exceed 90.

Once again, a what-if analysis of all the possible solutions to this problem would require examining and evaluating an unmanageably large number of scenarios – in this case, about 800,000. Optimization is necessary to find a good solution efficiently. A good starting point can probably be established by consulting with experienced managers in the insurance claims area, based on the expected demand of claims. We use OptQuest to optimize the staffing levels of this system. We run OptQuest for 100 iterations and 5 simulation runs per iteration. Table 2 shows a summary of the results, and Figure 4 shows the graph depicting the search of *OptQuest for Simul8* towards improving solutions.

Table 2: Summary Results

Simul8 Results										
Personal Claims Process - 20 Trials										
Solution	Claims Handler	Assessor (SC)	Assessor HQ	Senior Assessor	Doc. Control	Throughput (# claims)			Headcount	Avg. Cycle Time (min.)
						LSL	Avg.	USL		
1	9	17	17	15	16	1563.00	1568.00	1573.00	74	639.00
2	9	17	17	14	16	1559.00	1564.00	1570.00	73	658.00
3	8	17	16	15	16	1562.00	1567.00	1573.00	72	646.00
4	9	18	12	15	11	1611.00	1622.00	1633.00	65	503.00
5	9	18	11	15	11	1610.00	1621.00	1632.00	64	510.00

The performance graph shows the value of the objective (in this case Headcount) on the y -axis, and the number of iterations on the x -axis. The performance curve (indicated in green) shows only improving solutions.

Since some of the solutions obtained from our optimization are relatively close in terms of throughput and cycle time, an analyst may want to reevaluate a set of the best n solutions to assess the precision of the results. In Table 2, we present the best 5 solutions obtained from our OptQuest run, by conducting an experiment of 20 trials for each solution. The information can now be given to a process manager. The manager can analyze the trade-offs between headcount and throughput or cycle time, to decide which configuration best aligns with service levels and process goals.




Figure 4: Performance Graph for the Optimization of the Personal Claims Process

For example, we can see that solutions 1, 2 and 3 are statistically the same. Solutions 4 and 5 are significantly better than 1, 2 and 3 in terms of headcount, throughput and cycle time, so the analyst should pick one of these. Which one is better? We re-ran 60 trials for each of these two solutions, to obtain a 95% confidence interval for each of these measures. In both cases the confidence intervals overlap. In fact, the resulting confidence intervals for throughput are almost identical; the intervals for cycle time are also very close, with the expected cycle time for Solution 4 (503 minutes) about 1.4% lower than that for Solution 5 (510 minutes). The analyst should consider if the savings obtained from having one fewer Assessor justifies such a small difference in cycle time. If so, then solution 5 should be chosen.

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. The combination of two popular methods, *optimization* and *simulation*, has made it possible to overcome limitations of classical optimization approaches for dealing with uncertainty, where the goal is to find high-quality solutions that are feasible under as many scenarios as possible. Classical methods by themselves 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 as desired by the decision maker. In addition, as we demonstrate through illustrative examples in project portfolio selection,

emergency room operation and insurance claims staffing, modern optimization engines can effectively enforce requirements on one or more outputs from the simulation. Finally, simulation optimization produces results that can be conveyed and grasped in an intuitive manner, providing an especially useful tool for identifying improved business decisions under risk and uncertainty.

BIBLIOGRAPHY

- [1] W. F. Sharpe, "Capital asset prices: A theory of market equilibrium under conditions of risk," *Journal of Finance*, vol. 19, no. 3, pp. 425-442, 1964.
- [2] H. Markowitz, "Portfolio Selection," *Journal of Finance*, vol. 7, no. 1, pp. 77-91, 1952.
- [3] D. B. Hertz, "Risk Analysis in Capital Investment," *Harvard Business Review*, 1979.
- [4] J. April, M. Better, F. Glover and J. P. Kelly, "New advances and applications for marrying simulation and optimization," in *Proceedings of the 2004 Winter Simulation Conference*, 2004.