



INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Spreadsheet Inventory Simulation and Optimization Models and Their Application in a National Pharmacy Chain

Qiang Liu, Xinhui Zhang, Yan Liu, Lebin Lin

To cite this article:

Qiang Liu, Xinhui Zhang, Yan Liu, Lebin Lin (2013) Spreadsheet Inventory Simulation and Optimization Models and Their Application in a National Pharmacy Chain. INFORMS Transactions on Education 14(1):13-25. <https://doi.org/10.1287/ited.2013.0114>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.


The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2013, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>



Spreadsheet Inventory Simulation and Optimization Models and Their Application in a National Pharmacy Chain

Qiang Liu

Department of Biomedical, Industrial and Human Factors Engineering, Wright State University, Dayton, Ohio 45435; and School of Electromechanical Engineering, Guangdong University of Technology, Guangzhou, China, liuqiang@gdut.edu.cn

Xinhui Zhang, Yan Liu, Lebin Lin

Department of Biomedical, Industrial and Human Factors Engineering, Wright State University, Dayton, Ohio 45435
{xinhui.zhang@wright.edu, yan.liu@wright.edu, lin.30@wright.edu}

This paper presents a spreadsheet model for the solution of inventory management problems that arise in pharmacy chain stores. Though there have been abundant spreadsheet models for inventory management, most of these models focus on the use of embedded functions to compute complicated formulas. The spreadsheet model presented in this paper, however, incorporates a simulation of the ordering process and an iterative procedure to search for near-optimal solutions, and is easy to understand by students and practitioners. The methodology has been successfully implemented in a national chain of pharmacies; and the spreadsheet model integrated into the curriculum of several engineering courses including inventory management, simulation, and optimization and it has provided an interactive environment for students to experience real-life inventory decision making. This paper is dedicated to the memory of Dr. Paul A. Jensen, whose widely adopted Excel add-ins have inspired our education and research in spreadsheet models and provided the foundation for our work with pharmacy inventory systems.

Key words: spreadsheet modeling; teaching inventory management; simulation optimization

History: Received: August 2012; accepted: January 2013.

1. Introduction

Spreadsheets play an important role for academia and practitioners in operations research because they offer an effective visual model development platform. As a consequence, they have been widely used in various areas such as demand forecasting, resource optimization, simulation, and inventory management. In the area of inventory management, spreadsheet models have been abundant for many years, yet most of these models focus on the use of embedded statistical functions, such as NORMDIST and NORMSINV, to compute complicated inventory formulas (Chopra and Meindl 2012). In practice, these formulas can be so complex and hard to understand that students and inventory managers fail to accept them (Silver et al. 1998). The use of simulation for inventory analysis is an important topic in the simulation literature. The purpose of using simulation is to optimize system parameters, such as reorder points and order-up-to levels in an inventory system. Nevertheless, there seems to be a myth that computational approaches

for inventory problems are prohibitively expensive (Zheng and Federgruen 1991), therefore, simulation models are seldom used in practice for managing inventory.

A simulation-optimization methodology for inventory management was recently developed and successfully implemented at a national chain of pharmacies. The system has reduced out-of-stock prescriptions by 1.6 million per year, increased revenues by \$80 million per year, and reduced inventory by \$120 million for the pharmacy stores. The system was selected as one of the six finalists for the 2013 Edelman Award for achievement in operations research and management science (see Zhang et al. 2014).

The essence of the pharmacy inventory system was captured in a spreadsheet simulation model. The spreadsheet model translates a stochastic inventory problem into a deterministic optimization problem and serves as an excellent teaching tool in several upper-level undergraduate and graduate engineering

courses such as inventory management, simulation, and optimization, and is the focus of this study. The effective use of this tool has provided students a deeper understanding of several topics that include data visualization, probability and distribution fitting, inventory models, simulation models, optimization techniques, and the design of heuristics. The Edelman competition and the spreadsheet model bring real-world problems into the classroom, enhance students' practical problem solving skills, and cultivate lifetime learning for both students and industry-based practitioners.

The rest of this paper is organized as follows. Section 2 reviews selected traditional inventory models and the gap between education and practice. Section 3 presents examples of empirical and fitted demand distributions. Section 4 focuses on the details of the spreadsheet simulation model and Excel Solvers. Section 5 explains a local search-based iterative algorithm developed around the spreadsheet model for the inventory simulation problem. Section 6 describes how the techniques and accompanying materials have been adapted for use in our inventory, simulation, optimization and modern heuristics courses. Concluding remarks are offered in §7.

2. Traditional Inventory Models

Inventory theory is one of the most thoroughly studied areas in operations research, as demonstrated in the vast array of academic journal publications on the topic. In practice, inventory management has been recognized as one of the most important functions at a company, primarily because it has such a strong impact on its financial performance.

2.1. Traditional Analytical Spreadsheet Models for Inventory Management

Inventory management courses usually start with the economic order quantity (EOQ) model, followed by inventory policies and safety stock for continuous or periodic review systems. In the derivation of inventory policies, approximations to demand, such as normal and Poisson distributions are used to obtain analytically tractable formulas. Spreadsheets have long been used to visualize the results and to facilitate the calculation of the formulas; for example, see the textbooks such as Silver et al. (1998), and Chopra and Meindl (2012).

Spreadsheets are also often used in practice to guide the implementation of inventory management procedures. In these applications, an ABC analysis is typically carried out first to classify items and to set service levels. This is then followed by the calculation of economic order quantities and base stock using approximations of normal or Poisson distributions to achieve the specified service levels. For examples of

such an approach, see Murphy and Yemen (1986), Kleutghen and McGee (1985).

These models, however, may not always be applicable in practical settings. It is not uncommon to see the existence of various shapes of demand distributions that differ from standard distributions such as the normal and Poisson, which are popular in textbook inventory models. Though inventory analysis based on other distributions, such as the lognormal and gamma, is possible, the complexity in developing analytically appropriate inventory models and sophisticated formulas is a daunting task. These traditional inventory models, however, often result in confusion among students and on occasion considerable resistance from practitioners. These drawbacks reflect precisely the skepticism and reluctance of top management in pharmacy stores to go forward with such inventory management systems.

2.2. Spreadsheet Simulation and Optimization Model

Combining inventory simulation and optimization is another approach to solving the inventory problem. Of course, inventory simulation is an important topic by itself as affirmed, for example, by Law and Kelton (2000) and Jensen and Bard (2003). However, estimating optimal inventory policies, such as the reorder point and order-up-to level quantities, is a challenging problem. Kleijnen and Wan (2007) presented the optimization of an (s, S) system and compared the optimality of the response surface technique, perturbation analysis, and the popular OptQuest method based on the Karush-Kuhn-Tucker conditions. Fu and Healy (1997) presented several computational studies, including gradient-based, retrospective, and hybrid approaches to optimizing the simulation of (s, S) systems. Zheng and Federgruen (1991) proposed several bounds on the optimal solution of s and S , and developed an iterative algorithm for the computation of optimal (s, S) values. Nevertheless, the myth persists that sophisticated models are computationally difficult.

2.3. Inventory Management Education

To aid students to comprehend abstract concepts behind inventory theory, various hands-on classroom exercises and interactive games have been designed, for example, the Beer Game (Sterman 1989, Simchi-Levi et al. 2003), the Inventory Control Exercise at Spiegel Grove (Drake and Mawhinney 2007), and an in-class competition of the multi-item newsvendor problem (Robb et al. 2010).

These games and exercises have created opportunities for students to experience life as supply chain decision makers and effectively enhanced students' learning experience (Drake and Mawhinney 2007).

Nevertheless, it is our experience that a common question raised by students is “how do we use these models in practice with thousands if not millions of items that could be of various cost structures and from various distributions.” The use of real-life data in rigorous yet visual inventory studies under various demand distributions would greatly enhance student’s confidence to carry out and promote inventory management in practice.

The simulation-optimization methodology for inventory management for the national pharmacy chain has several characteristics that fulfill this need: (a) it is based on the simulation of inventory systems; the process of constructing the inventory model helps students understand how inventory ordering policies are implemented and the process can be visualized, (b) it is versatile and is not limited to specific demand distributions; it allows the use of empirical distributions while avoiding complicated formulas and the resistance that accompanies them, and (c) it connects the analytic models (if desired, analytic solutions can serve as a seed to start the optimization process) and offers an alternative approach to solving real-world problems.

The material discussed in this paper has been adopted in (a) inventory management—an upper-level undergraduate and graduate engineering course for demand distribution fitting, inventory simulation and optimization, and visual tools for practical inventory management; (b) simulation—an upper-level undergraduate and graduate engineering course for spreadsheet-based simulation optimization; (c) deterministic optimization techniques—an upper-level undergraduate and graduate engineering course in which case studies for spreadsheet modeling with the Excel Solver play a prominent role; (d) modern heuristics—a graduate engineering course, in which neighborhood search examples are drawn from the pharmacy study; and (e) successful application of operations research—a graduate level seminar course based on industrial case studies. The material contains more than one year of transaction history of 82 drugs, provides a rich set of distributions, provides an interactive environment for students to experience real-life inventory decision making, and was much appreciated by the students.

In the following sections, we present the details of the spreadsheet simulation model for the pharmacy inventory problem. Despite the fact that this model is developed for a specific pharmacy chain, many of its elements can be applied to other inventory systems.

3. Demand Distributions in Pharmacy Stores

The pharmacy chain that sponsored this research operates more than 1,950 stores in the United States

with thousands of items (drugs) at each store. To start any inventory project, for students, it is recommended to investigate the demand pattern in the specific environment. Such investigations are necessary for both simulation and inventory management courses.

3.1. Multimodal Demand Distributions for Customer Orders


Pharmacy demand exhibits several characteristics that differentiate it from demand in other industries; for example, multiple streams of demand of specific order sizes create multiple peaks in the demand distributions. As a result, many of these distributions are multimodal and do not fit the unimodal statistical distributions seen in traditional inventory models. For example, Figure 1 illustrates historical demand (top panel) and empirical and fitted lead time demand distributions (bottom panel) of such a drug.

In this simple example, the demand mainly falls into a few discrete values—the basic dose for this drug is one pill per day; thus, a 30-day supply is 30 pills, a 60-day supply is 60 pills, and a 90-day supply is 90 pills. It is clear that the demand (top panel) reflects the occurrence of multiple independent streams of 30-day, 60-day, and 90-day demand from customers, though an occurrence of a 90-unit demand could be either one 90-day supply for one customer, or some combination of 30-day and 60-day supplies, or three independent 30-day supplies for three different customers. In a pharmacy, the demand for the vast majority of the drugs is highly variable, intermittent, and irregular; and if the 90-day demand and 30-day demand exceed the 60-day demand, as is the case for this drug, two peaks in the demand distribution result, as shown in the bottom portion of the figure; thus, an accurate model requires a multimodal distribution.

3.2. Demand Distribution for Drugs at Pharmacy Stores

Though not all empirical demand distributions are multimodal, it is our observation that much of the lead-time demand cannot be modeled using standard theoretical distributions. For reference, Figure 2 shows several commonly seen demand distributions observed in a typical store.

The lead-time demand, as suggested through goodness of fit test, can be approximated by several shapes, such as normal, Poisson, or gamma, and the distribution could be multimodal, and thus may not fit neatly into the standard unimodal distributions seen in traditional inventory models. Though in certain cases fitting an individual unimodal distribution would be a viable approach (see [Nenes et al. 2010](#) for examples) such an approach requires subjective judgment to determine which theoretical distribution

Figure 1 Historical Data Show Multiple Streams of Independent 30-, 60-, and 90-Day Demand, Shown in the Top Panel of the Figure

Note. These multiple streams of demand imply multiple peaks in the underlying lead-time demand histogram, shown in the bottom panel of the figure, suggesting that a unimodal distribution is not appropriate for modeling the lead-time demand distribution.

should be used, and the resulting bundle of mathematical models is difficult for management to understand. Under these circumstances, simulation-based models, such as the one proposed in this study, provide a visual, intuitive, and powerful approach to solving inventory problems.

4. The Spreadsheet Simulation Model for Pharmacy Inventory Management

In this section, we present our spreadsheet simulation model developed to mimic the pharmacy inventory ordering process, followed by optimization using

the Excel Solver to find “optimal” inventory policies. This material is well suited for simulation, inventory management, and optimization courses designed for upper-level undergraduate or graduate engineering students.

4.1. Inventory Management at Pharmacy Stores

Drug orders as well as nongrocery items are delivered to pharmacy stores by truck following weekly schedules, e.g., Monday, Wednesday, and Friday. The pharmacy stores use a periodic review (s, S) system, where s is the reorder point and S is the order-up-to level; the lead time is the number of days between deliveries.

Figure 2 Pharmacy Drugs Shows Various Distributions and Could Be Multimodal





Figure 2 Continued



At the beginning of each period, a pharmacist checks the inventory position of the drugs. If it is at or below s , the pharmacist will create an order to raise the inventory level to S . If it is above s , on the other hand, nothing will be done until at least the next review period. An order is usually rounded up to a multiple of a certain bottle size, such as 100 units per bottle, and arrives right before the next delivery period. Once an order is placed, an estimated cost is incurred for the warehouse picking, transportation, and store put-away activities.

When a prescription order (demand) arrives, it is satisfied immediately if it can be satisfied by in-stock inventory. If the demand exceeds the inventory, it is assumed that the demand is unfulfilled and a shortage cost is incurred per stock-out. The shortage cost includes the profit loss of the prescription and an estimated percentage loss of all future purchases from the affected customer. The pharmacy chain is cautious

on its inventory investment and sets a high inventory holding cost to drive down its inventory, eliminate potential obsolescence, and release the capital that would otherwise be tied up in inventory. As such, the fundamental inventory simulation-optimization problem is to find the optimal inventory policies which, in this study, are the values of s and S , so as to minimize the total inventory holding, ordering, and stock-out costs.

4.2. The Spreadsheet Simulation Model

The spreadsheet simulation model (supplemental file PharmacySimuOptConcise.xlsx; available as supplemental material at <http://dx.doi.org/10.1287/ited.2013.0114>) takes as inputs the historic demand from a store (the “TransLog” worksheet), simulates the ordering process of a particular drug for a given (s, S) policy, and outputs the performance of the inventory system.

Figure 3 Spreadsheet Model Shows the Performance of Various Inventory Policies and Its Optimization Through Excel Solve




Figure 3 illustrates the spreadsheet simulation model, which consists of four sections: the inventory policy or decision variable section, the process simulation section, the graphic section, and the results section.

The decision variable section (D2:F5). Cell B2 represents the drug to be optimized by its national drug code (drugID), cell B3 the number of prescriptions of the drug, cell B5 the package size of each order, cell D3 the value of the reorder point (s), and cell F3 the value of the order-up-to level (S). Cells D3 and F3 are the decision variables, quantities for (s , S), respectively.

The process simulation section (A6:K32). Column A is the query key used to aggregate the demand of prescriptions for the drug on any particular day (listed in column B). The aggregate demand is shown in column C. Column D represents the days of a week, which is used to determine whether or not this is a review period (shown in column E, where 1 represents a review period and 0 otherwise). In the example, the review periods are Monday, Wednesday, and Friday. Column F represents the current inventory on hand, and column G the inventory positions. Column H represents the decision to order (coded as 1) or not. If the inventory position is less than or equal to s and the day is a review period, then an order, rounded to a multiple of the package size, 10 as defined in cell B2, is issued to reach S . This order quantity is shown in column I. The dates and quantities of order arrivals, shown in column J, are calculated based on the lead time of the drug.

Columns K and L keep the ending inventory and the estimated number of out-of-stocks of the drug on each day, respectively.

The graphic section (N7:R32). The drug's average inventory at the end of each day and its aggregate demand every day are shown in this section. The plot of the end of the day inventory gives a vivid picture of the inventory policies defined by (s , S). If s is set too high, inventory levels will be above zero; if s is set too low, stock-outs will occur.

The results section (J2:N3). The total number of orders, the average inventory, and the total number of out-of-stock prescriptions are calculated in cells J3, K3, and L3, respectively. These measures are translated into an objective value shown in cell N3.

The spreadsheet model can be easily connected to an enterprise information system to retrieve the transactions for a store, which are stored in the "TransLog" worksheet. By changing the national drug code, (cell B2), the demand for that drug is automatically retrieved from the datasheet through the concatenation of drug-year-month-day as the query key to generate the fixed demand.

For a given drug to be optimized, by fixing this demand, the spreadsheet simulation model mimics the inventory process for the given (s , S) policy, and translates the policy's performance into a single value defined as the total ordering, holding, and stock-out costs. As such, we have translated a stochastic inventory problem into a deterministic optimization problem where various techniques can be used to find the optimal solutions.

4.3. Use of Standard Excel Solver for the Inventory Simulation-Optimization System

One of the major benefits of MS Excel is the built-in Solver, which allows business users to solve various optimization problems. Based on the simulation model described above, by fixing the demand, Solver can be conveniently used to solve the deterministic optimization problem, transformed from the stochastic inventory problem, to find the optimal or near-optimal inventory policy, defined by variables (s , S).

There are three optimization methods built into the Excel Solver: simplex LP, GRG nonlinear, and evolutionary. The simplex LP is designed for linear programs, the GRG nonlinear engine for smooth nonlinear programs, and the evolutionary engine for non-smooth problems. Here, we again use the drug shown in Figure 3 as an example to demonstrate the performance of the Excel Solvers and justify the motivation of our algorithm. Here cell N3 holds the objective to be minimized, and D3 and F3 hold the decision variables, s and S , respectively.

The inventory simulation and optimization model is highly nonlinear and both the GRG nonlinear solver and evolution algorithm can be used. Our computational experience showed that (a) If the GRG nonlinear solver is chosen, the quality of the solutions obtained often depended on the initial solution provided; (b) if the evolutionary solver is chosen, typically, the solution quality is higher, but convergence is slower. In any case, Excel Solver is able to find near-optimal solutions to the example problem, as shown in Figure 4 (in this case, evolutionary algorithm, was used). In this solution, the total orders are 53, the average inventory is 163 units, the stock-out is 0, and the objective value is 11.96.

In concluding this section, we would like to mention that the spreadsheet model is easy to construct by undergraduate students (though certain knowledge of spreadsheet modeling is required). It strengthens students' appreciation of the ordering process and provides a clear visual of the seesaw inventory pattern. Moreover, Solver produces near-optimal solutions that are understood by all students who get hands-on experience in solving real-world inventory problems. As a final point, it should be noted that the performance of the default Excel Solver was not wholly satisfactory. Consequently, to obtain robust near-optimal solutions quickly, customized algorithms were designed.

5. Local Search Heuristic Design for the Inventory Simulation-Optimization Problem

Modern heuristics are becoming an integral part of operations research curricula for graduate study because of their capability to solve difficult optimization problems, for example, the well-known vehicle routing problem. In this section, we describe the development of a local search heuristic for the solution of inventory simulation-optimization problems. This material could be used in a heuristic design course for graduate students.

5.1. Motivation for the Heuristics

The spreadsheet simulation provides a rich source of insights from which we have devised heuristic procedures to “move” or “adjust” the (s , S) values in the search for near-optimal inventory policies. The heuristic mainly alternates between two phases or procedures: the first attempts to find the reorder point under specific order sizes to achieve a balance of inventory and shortage costs; the second is aimed at balancing the ordering and inventory costs.

These procedures are accomplished through several moves. Specifically, procedure 1 is accomplished through move 1 by increasing or decreasing (s , S) both by an equal amount; and procedure 2 is accomplished through move 2 by increasing s while keeping S unchanged to decrease the order size Q , or move 3 by increase S while s unchanged to increase the order size Q . Though the moves or procedures are similar to those used in the literature, the most salient feature of our heuristic is the selection of step size to achieve quick convergence.


To illustrate, we present an example of how these moves work for the inventory problem outlined in Figure 3. For this particular problem, the demand is mostly 30, 60, or 90 units, with occasional demand of 10, 15, or 45 units. The package size is assumed to be 10 units per bottle and an order has to be a multiple of this package size.

Starting Solution (180, 220). For simplicity, we start the algorithm with the maximum demand in the order period; in this case, this value turns out to be 180. The economic order quantity is calculated to be 40, which gives the solution of (180, 220) with an objective value of 25.17. The calculation of starting solution is shown in cell P3 (EOQ), Q3 (s), and R3 (S)

Figure 4 Excel Solver Was Able to Find Near Optimal Solution, Yet Its Computational Performance Calls for Customized Algorithms

Decision variables				Performance metrics			Objectives
Reorder point		Order-up-to level		Total orders	Average inventory	No. of OOS	Objective
165		235		53	163.30	0	11.96

Figure 5 Simulation of (180, 220) Suggests Either an Increase of s by 20 Units to Increase Negative Inventory from -20 to 0 or a Decrease of s by 10 Units to Bring the Next Positive Inventory from 10 to 0 in Search of Local Optimum Under $Q = 40$



in the supplemental spreadsheet file. Figure 5 shows how ending inventory (top panel) and demand (bottom panel) evolve over time for $(s, S) = (180, 220)$.


Let us start with strategy 1 and move 1. As we can see, the maximum negative ending inventory is -20 and the minimum positive inventory is 10. If we are to increase s , it is intuitive to increase it by 20 to bring the next negative inventory to 0. If we are to decrease s , it is intuitive to decrease it by 10 units so as to bring the next positive inventory to 0. Increasing the reorder point, from 180 to 200, in this case, leads to a better objective and is adopted. By keeping Q at 40, we can move to the next inventory policy, given by $(200, 240)$ with an objective value of 12.97, for which the ending inventory is shown in Figure 6.

Solution 2. Inventory policy $(200, 240)$ exhibits no excess inventory or stock-out and also has no negative inventory. Inventory policy $(200, 240)$ is the local optimal inventory policy under $Q = 40$, as no move 1 will lead to better solutions. This concludes the search in procedure 1.


Once procedure 1 reaches local optimum, a change of order size Q is necessary and this is performed in procedure 2 where we can either increase or decrease Q . Silver et al. (1998) state that when s and S (or Q) are determined simultaneously, Q is always larger than the EOQ; therefore, to start, we decided to increase the size of Q .

Notice that in the ending inventory plot for solution 2, $(200, 240)$, the minimum positive inventory

Figure 6 Simulation of (200, 240) with the Objective Value of 12.97 Under $Q = 40$



Note. The minimum inventory is 30, which suggests a potential change in order size of 30 units will be required to find a better solution.

Figure 7 Simulation of (170, 240) Suggests a Parallel Reduction in Both s and S by 5 (Move 1)


Note. This is because all ending inventory are above 0.

level is 30. Here, if we are to increase the size of Q , there are two options available. The first option is to apply move 2, which decreases s by 30 while keeping S the same at 240; here we would move from (200, 240) to (170, 240). The second option is to apply move 3, which increases S by 30 units while keeping s the same at 200; here we would move from (200, 240) to (200, 270). By increasing the order size Q , the policy is likely to increase the inventory holding cost. Here, the first option is chosen so as to offset this inventory increase, so we end up at (170, 240). Nevertheless, the second option, if applied, will lead to the same local optimal solution under the same order size ($Q = 70$) and is not shown here. Similar moves exist for the

case where Q is decreased and is not explained here in detail.

Solution 3. Inventory policy of (170, 240) has an objective value of 12.31, for which the ending inventory is shown in Figure 7. Notice that the ending inventory was 5, well above 0, suggesting a parallel reduction in both s and S , move 1, which leads to the solution (165, 235).

Solution 4. Inventory policy (165, 235) has an objective value of 11.96, for which the ending inventory is shown in Figure 8. This solution was confirmed to be best solution in the discrete solution space through an exhaustive enumeration of all possible (s , S) combinations.

Figure 8 Simulation of (165, 235) Was Locally Optimal with Respect to Q , and Later Confirmed to Be the Global Optimum

5.2. The Algorithm

In the presentation of the algorithm, let f be the objective function value and superscript $'$ the neighboring solution. Let Q^+ , Q^- be an upper and a lower bounds for Q , respectively, and ΔQ be the direction of change for Q , i.e., $\Delta Q = +1$ represents an increase of Q ; $\Delta Q = -1$, a decreases of Q . Similarly, let Δ_s be the change of direction for s , i.e., $\Delta_s = +1$ represents an increase of s ; $\Delta_s = -1$, an decreases of s . Let B be the set of order quantities, and let P be the set of inventory policies encountered in the search process. P consists of the three-tuple (s, S, f) .

Our iterative algorithm for finding near-optimal (s, S) policy is the following:

Step 1: Initialization. Start the algorithm with $s = s'$, the reorder point, as the maximum demand in a period, $Q = \text{EOQ}$ (the economic order quantity), $S = S' = s + Q$, f , the current objective is set to $+\infty$, set $\Delta Q = +1$ to increase Q and $\Delta_s = 1$, to increase the reorder point. Push Q into set B .

Step 2: Simulation. Perform simulation under (s', S') , denote the simulation objective value as f' , push policy (s, S, f) into set P . Denote $I(t)$ as the ending inventory of each period, and let $I^+ = \min_{t: I(t) > 0} \{I(t)\}$ be the minimum positive inventory, $I^- = \min_{t: I(t) < 0} \{|I(t)|\}$ be the absolute value of the maximum negative ending inventory.

Step 3: Increase or Decrease Reorder Point. If $f' < f$, continue search as follows. If $\Delta_s = +1$, set $s' = s' + I^-$ and $S' = s' + Q$, in an effort to raise the smallest negative inventory to zero. Else $\Delta_s = -1$. Set $s' = s' - I^+$ and $S' = s' + Q$ in an effort to lower the next positive inventory to zero, go to Step 2. Otherwise, $f' > f$, local optimal under Q is found, go to Step 4.

Step 4: Increase or Decrease Q ? Let $Q\Delta$ be the minimum of I^+ and I^- , if $\Delta Q = +1$, then $Q = Q + Q\Delta$; otherwise, $Q = Q - Q\Delta$. If no Q exist inside (Q^-, Q^+) , set $\Delta Q = -\Delta Q$, reverse search directions for order size Q . If Q exists in set B already, go to Step 6; otherwise, push Q into set B , go to Step 5.

Step 5: Perform moves 2 or 3 to a new (s', S') if $\Delta Q = +1$, then move to $(s' = s' - Q\Delta, S = S)$ and set $\Delta_s = 1$. If $\Delta Q = -1$, then move to $(s' = s, S' = s' + Q)$, set $\Delta_s = -1$. Go to Step 2.

Step 6: Termination. Output the best solution in set P

In Step 1, we start with an order size equal to EOQ. Silver et al. (1998) state that in simultaneous determination of s and S , (or Q), Q is always larger than EOQ. As such, we set $\Delta Q = +1$ to increase EOQ. The reorder point can be set in different ways; for example, the maximal demand in lead-time, or with any analytical approximation solutions.

In Steps 2 and 3, we perform simulation with given values of s and S . Based on the ending inventory resulting from the simulation, we adjust both

the reorder and order-up-to quantities by the same number so as to keep the order size unchanged. Specifically, if the current solution is improving, then we continue the search until a local optimum under a fixed order size is obtained. In making these adjustments, if we increase the reorder point, then the algorithm increases the reorder by I^- to bring the next negative inventory to zero (reducing stock-out cost); if we decrease the reorder point, then the algorithm decreases the reorder point by I^+ to bring the next positive inventory to zero (reducing inventory cost).

In Steps 4 and 5, we perform an adaptive search with respect to the order size, Q . Rather than enumerate all possible breakpoints as proposed by Fu and Healy (1997), which could be time consuming, we increase or decrease the order size based on the ending inventory level derived from the simulation. Finally, we stop if no improvement has been made or Q is outside the predefined bounds and output the final best solution encountered in the search process in Step 6.

Computational Summary. Our computational experiments showed that for most of the drugs stocked by the sponsoring company, the algorithm converged in only 10 to 20 milliseconds on an Intel i7 desktop. See Zhang et al. (2014) for details.

6. Suggested Use and Classroom Results

The simulation-optimization method for inventory management adopted by the sponsoring pharmacy chain has been used by the corresponding author in courses such as (a) inventory management, as visual tools for solving practical inventory problems; (b) simulation, as material for spreadsheet-based simulation optimization; (c) optimization techniques, as case studies for spreadsheet optimization; (d) modern heuristics, as material for neighborhood search; and (e) successful applications of operations and research, as case studies in industry.

Each of the §§3, 4, and 5 centers on a specific topic and can be reorganized to fit the specific course. For example, §§3 and 4 can be combined and used in inventory management or simulation or an optimization course, and §§4 and 5 can be combined and used in a heuristic course. The following questions in conjunction to the sections selected provide a basis for further investigation to suit an instructor's need.

Section 3: Distribution fitting. (a) What is the best fitting distribution(s) for the demand distribution(s) from industry data? (b) Why do the normal distributions appear in fast-moving drugs whereas Poisson distributions appear in slow-moving drugs? (c) What causes demand in a pharmacy store to follow a multimodal distribution? (d) What are the errors induced

when a wrong distribution is selected for an inventory model? (e) In what cases can I use normal or Poisson distributions? (f) How is a distribution selected that best fits a particular demand?

Section 4: Spreadsheet simulation and optimization. (a) construct the spreadsheet model from the data sheet; (b) run the simulation and use Solver to find the near optimal solutions; (c) determine how the simulation model can be changed from a periodic to a continuous review system and how the solution will be affected; (d) given a continuous review system, determine if we change the demand to a normal or Poisson distribution, and compare the corresponding analytic solutions with the simulation results; (e) perform a grid search on (s, S) , table the results, evaluate the performance of Solver under various options and why starting solutions affect the convergence of the GRG algorithm, yet not the evolutionary algorithm; how can the optimization algorithm be speeded up when additional constraints such as bounds being included; and (f) determine the benefits provided by the simulation-optimization model and will different forms of demand distributions affect the simulation-optimization models?

Section 5: Heuristic design. (a) Determine the local moves used and how they can be controlled; (b) manually perform these moves to gain insight into the selection of moves and the step size of a move; (c) consider how an approximate solution can be used to speed up the solution process; (d) implement the algorithm using a language of your choice.

These considerations extend the models beyond the pharmacy application. For example, by changing the “Demand” column (column C in Figure 3) to a normal distribution, say $\text{NORMINV}(\text{RAND}(), 100, 25)$ where 10 is the mean, 25 is the standard deviation, and by changing the “order or not” column (column E in Figure 3) to 1 for a continuous review system, the spreadsheet model can serve as a visual tool to evaluate the traditional inventory policies. If desired, the optimization can be used to jointly optimize both the reorder-point and order-up-to quantities, and compare the solution to iterative procedures for solving this problem; see Silver et al. (1998) for such procedures. With these issues in mind, we now outline how the material can be inserted into our operations research curriculum.

Inventory management. The focus is inventory optimization through simulation. The students are asked to fit a demand distribution, explain the reasons behind the existence of multimodal distributions, construct the simulation model, and run the simulation; to evaluate various inventory policies, compare them with analytic solution obtained by using, say, a normal demand—by changing “demand” column (column C in Figure 1) with normal distributions,

visualize solutions; and to compare periodic and continuous systems by changing the values in the “ordering period” column (column E in Figure 1) to 1.

Simulation. The focus is on a combination of simulation and optimization through Excel. The students are asked to fit distributions, construct the simulation model, run the simulation, and run the Excel Solver. In the case of random demand—by changing the “demand” column (column C in Figure 1) to reflect normal distributions, they are asked to run multiple scenarios and perform statistical analysis of the output.

Modern heuristic. The focus is on algorithm design so the simulation model is provided to the students. The students, however, are asked to manually evaluate the moves based on simulation output, understand the procedures or meta-strategies to guide the moves, provide a pseudocode, and then write a computer code using a language of their choice.

This material has received positive feedback from students and helped them to better understand distribution fitting, inventory models, and the ordering process, simulation construction, and optimization techniques. For the inventory class, it should be mentioned that the material is not intended to replace the current analytic approaches that are traditionally taught; rather, the spreadsheet model can serve as a tool for students to understand inventory models—the process of constructing the inventory model helps students understand how inventory is ordered and how the results can be visualized, gives students confidence on the use of these models, and is mostly appreciated by the students.

However, it is our experience that for undergraduate students to successfully construct the spreadsheet simulation model, it is necessary to include basic materials on spreadsheet modeling; a programming language is a prerequisite for the modern heuristics course for graduate students to be able to write a pseudo-code or a computer code for algorithmic implementation. Such prerequisites, however, should already exist for any course in heuristics.

7. Concluding Remarks and Special Acknowledgement to Dr. Jensen

The material described in this paper along with the material developed for the Edelman competition offers intuitive tools for helping students to better understand spreadsheet modeling, inventory control, simulation and optimization, heuristic design, and operations research in general. Using a case study format, it bridges all of these topics while enhancing the skills needed by analysts to solve practical inventory problems. It is our belief that the innovative approach

described herein can be applied in many retail contexts and so strengthens lifelong learning for both students and practitioners alike.

This paper is dedicated to the memory of Dr. Paul A. Jensen. While the corresponding author, Xinhui Zhang, was a doctoral student at the University of Texas at Austin he benefited tremendously from Dr. Jensen's enthusiasm toward operations research applications, and he remains grateful for the guidance Dr. Jensen provided as a doctoral committee member. Since graduating he has used Dr. Jensen's Excel add-ins for years in his teaching of optimization, queueing, and simulation, and he grew to understand their enormous potential. In fact, they provided the foundations for the design of the spreadsheet models described here.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/ited.2013.0114>.

References

- Chopra S, Meindl P (2012) *Supply Chain Management: Strategy, Planning and Options* (Prentice Hall, Upper Saddle River, NJ).
- Drake MJ, Mawhinney JR (2007) Inventory control at Spiegel Grove. *INFORMS Trans. Ed.* 8(1):19–24.
- Fu MC, Healy KJ (1997) Techniques for optimization via simulation: An experimental study on an (s, S) inventory system. *IIE Trans.* 29(3):191–199.
- Jensen PA, Bard JF (2003) *Operations Research: Models and Methods* (John Wiley & Sons, Hoboken, NJ).
- Kleijnen JP, Wan J (2007) Optimization of simulated systems: OptQuest and alternatives. *Simulation Modelling Practice and Theory* 15(3):354–362.
- Kleutghen PP, McGee JC (1985) Development and implementation of an integrated inventory management program at Pfizer pharmaceuticals. *Interfaces* 15(1):69–87.
- Law AM, Kelton WD (2000) *Simulation Modeling and Analysis* (McGraw-Hill Education, Berkshire, UK).
- Murphy J, Yemen S (1986) Computer-assisted inventory control utilizing ABC inventory analysis and EOQ in a hospital pharmacy. *The Canadian J. Hospital Pharmacy* 39(6):159–163.
- Nenes G, Panagiotidou S, Tagaras G (2010) Inventory management of multiple items with irregular demand: A case study. *Eur. J. Oper. Res.* 205(2):313–324.
- Robb DJ, Johnson ME, Silver EA (2010) An in-class competition introducing inventory management cepts. *INFORMS Trans. Ed.* 10(3):122–125.
- Silver EA, Pyke DF, Peterson R (1998) *Inventory Management and Production Planning and Scheduling* (John Wiley & Sons, New York).
- Simchi-Levi D, Kaminsky P, Simchi-Levi E (2003) *Designing and Managing the Supply Chain*, Second ed. (Irwin/McGraw-Hill, Boston).
- Sterman JD (1989) Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Sci.* 35(3):321–339.
- Zhang X, Meiser D, Liu Y, Bonner B, Lin L (2014) Simulation and optimization improves pharmacy inventory management at Kroger Corp. *Interfaces*. Forthcoming.
- Zheng Y, Federgruen A (1991) Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Oper. Res.* 39(4):654–665.