Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

IDENTIFYING DEMAND SOURCES THAT MINIMIZE RISK FOR A SELECTIVE NEWSVENDOR

Kevin M. Taaffe
Deepak Tirumalasetty

Department of Industrial Engineering
Clemson University, 110 Freeman Hall
Clemson, SC 29634-0920, U.S.A

ABSTRACT

Consider a firm that offers a product during a single selling season. The firm has the flexibility of choosing which demand sources to serve, but these decisions must be made prior to knowing the actual demand that will materialize in each market. Moreover, we assume the firm operates on a tight budget and cannot afford to record several successive financial losses spanning consecutive periods. In this case, it is likely that their objective is not only to maximize expected profit, but to minimize the variance from that goal. We provide insights into the tradeoff between expected profit, expected revenue, and demand uncertainty. Finally, we present a solution approach, via simulation, to determine the best set of markets to pursue and the associated order quantity when the firm's objective is to minimize the probability of receiving a profit below a critical threshold value.

1 INTRODUCTION

As product lives continue to decrease with technological advances and fashion trends, and the efficiency of manufacturing processes offer less room for improvement, a supplier or manufacturing firm is constantly trying to identify other ways to improve profitability. In the classic newsvendor problem, the firm seeks an optimal procurement policy for a product with random demand during a single selling season. There is extensive literature on this topic, and we refer the reader to Porteus (1990), Tsay, Nahmias, and Agrawal (1999), Cachon (1999), and Petruzzi and Dada (1999) for reviews and research in this area.

If the firm can obtain unique revenues in each demand source (or market), then the problem becomes one of simultaneously selecting the most desirable *markets* as well as determining the appropriate total *order quantity* before demand is actually realized. Recent research has offered profit maximizing models that provide integrated demand selection and ordering decisions for this so-called "selective newsvendor" problem (SNP). Forms of the SNP have been studied recently by Carr and Lovejoy (2000), Petruzzi and Monahan (2003), Taaffe, Geunes, and Romeijn (2005), Taaffe and Romeijn (2005), and Taaffe, Geunes, and Romeijn (2004).

In both categories of the aforementioned problems, the typical objective is to maximize expected profit or minimize expected cost, which would be appropriate for a risk-neutral firm. However, not all (in fact, very few) firms have the luxury of operating in a risk-neutral environment (see Schweitzer and Cachon (2000)). The actual profit (or loss) may be quite different than expected profit for a particular selling season, and many firms could be more concerned with this variability. Therefore, we consider a firm that cannot afford successive losses or negligible profits spanning several selling seasons. For such a firm, we will evaluate two risk models. In one situation, we still assume that the firm's objective is to maximize expected profit, but now it will also require that a given percentage of outcomes or profit realizations must achieve some minimum desired profit level. In a second situation, while the firm's desire may be to maximize expected profit, their objective will be to minimize the number of outcomes that could occur below their budgeted profit level.

Various aspects of risk aversion in newsvendor problems have been considered in past work. Lau (1980) is the first paper to directly study the effect that risk has on the newsvendor problem. The paper considers two objectives, maximizing expected utility, and maximizing the probability of achieving a budgeted profit, which is quite similar to the focus of our work. However, we have the added complexity of simultaneously selecting the most attractive markets while determining the appropriately-sized order quantity. Lau (1980) depicts two demand points beyond which the firm will no longer achieve the desired profit level, and then solves for the Q that maximizes the probability that the profit level will be achieved. The paper concludes that analytical solutions can be obtained if the underlying demand distribution is normal or exponential. This approach works for a standard newsvendor when there is only one demand distribution for which all demands generate the same per-unit revenues. Applying this methodology to our problem breaks down due to our unique revenues in individual markets.

Eeckhoudt, Gollier, and Schlesinger (1995) also studies a risk averse newsvendor for which any demands not met by the original order can be satisfied through a high-cost local supplier. This paper also concluded that the optimal risk-averse order quantity is less than the amount ordered in the expected value solution. More recently, Collins (2004) offers some results that counter these previous papers.

Collins (2004) conjectures that there is a class of problems for which the risk averse and expected value solutions are identical, that there are many problems for which the expected solution provides a good approximation to the risk averse solution, and that in most problems in practice, the risk averse solution would actually be to order *more* than the expected value solution.

Finally, the reader can turn to Chen et al. (2004) and Van Mieghem (2004) for additional risk aversion research.

Our work in this paper will be to evaluate how a selective newsvendor will integrate risk into its demand selection and ordering policy. While we maintain some similar assumptions to those in Lau (1980) and Eeckhoudt et al. (1995), we also have the added complexity of market selection, which can potentially result in different procurement policies. In Section 2, we introduce the general profit equation for the selective newsvendor problem and discuss the form of the distribution for profit. Then, in Section 3, we introduce two alternate demand selection formulations, both of which include some focus on risk. We also provide our heuristic solution approaches to each problem. In Section 4, present computational tests and findings for each model. Finally, we summarize our findings in Section 5.

2 QUANTIFYING PROFIT FOR THE DEMAND SOURCE SELECTION PROBLEM

For a complete treatment of the selective newsvendor with no risk aversion, please refer to Taaffe, Geunes, and Romeijn (2004). We begin by defining c as the per-unit cost of obtaining or procuring the product to be sold. The product can be sold in market i at a per-unit price of r_i. If realized demand is less than the quantity ordered, the firm can salvage each remaining unit for a value of v. If demand exceeds the order quantity, there is a shortage cost of e per unit. However, we assume that the demand is still met through expediting via a local supplier or single-period backlogging whereby a second order can be placed with the firm's regular supplier. In either case, the unit cost is still e.

Recall that, in the selective newsvendor framework, the firm must decide its market selections prior to placing the order for Q units. Let $y_i = 1$ if the firm decides to satisfy demand in market i, and 0 otherwise. We present the following expression for the total realized profit, based on the order quantity, market selection decisions, and realized demand.

$$G(Q, y) = \begin{cases} \sum_{i=1}^{n} (r_i D_i - S_i) y_i - cQ + v(Q - \sum_{i=1}^{n} D_i y_i) \\ \qquad\qquad \text{if } Q > \sum_{i=1}^{n} D_i y_i \\ \sum_{i=1}^{n} (r_i D_i - S_i) y_i - cQ - e(\sum_{i=1}^{n} D_i y_i - Q) \\ \qquad\qquad \text{if } Q \le \sum_{i=1}^{n} D_i y_i \end{cases}.$$

Given a binary vector of market selection variables y, let $D^y = \sum_{i=1}^n D_i y_i$ denote the total demand of the selected markets, and denote its pdf by f_y and its cdf by F_y. It is easy to see that the total selected demand has mean $E(D^y) = \sum_{i=1}^n \mu_i y_i$ and variance $\mathrm{Var}(D^y) = \sum_{i=1}^n \sigma_i^2 y_i$. We can then express the firm's expected profit as a function $\hat{G}(Q, y)$ of the order quantity Q and the binary vector of market selection variables y:

$$\hat{G}(Q, y) = \sum_{i=1}^{n} (r_i \mu_i - S_i) y_i - cQ + v \int_0^Q (Q - x) f_y(x) \mathrm{d}x - e \int_Q^\infty (x - Q) f_y(x) \mathrm{d}x.$$

For a given vector y, the expected profit function $G(Q, y)$ is concave, and maximizing the expected profit is equivalent to minimizing the cost in the associated newsvendor problem. We can then derive the optimal order quantity based on the critical fractile of total demand, or $\rho = \frac{e-c}{e-v}$. Thus, the optimal order quantity will be $Q_y^* = F_y^{-1}(\rho)$ (see Taaffe et al. (2004)). When the firm faces normally distributed demand in each market, the total demand satisfied (i.e., $D^y = \sum_{i=1}^n D_i y_i$) will always follow a normal distribution, which allows the use of the standard normal loss function in reducing the expected profit equation to

$$\hat{G}(Q, y) = \sum_{i=1}^{n} \bar{r}_i y_i - K(c, v, e) \sqrt{\sum_{i=1}^{n} \sigma_i^2 y_i}, \qquad (1)$$

where $\bar{r}_i = ((r_i - c)\mu_i - S_i)$, and $K(c, v, e) = \{(c - v)z(\rho) + (e - v)L(z(\rho))\}$. Thus, the expected profit equation depends solely on market selection variables, and the optimal order quantity is simply a function of y, given by $Q_y^* = \sum_{i=1}^n \mu_i y_i + z(\rho)\sqrt{\sum_{i=1}^n \sigma_i y_i}$.

We make a key observation here. We previously stated that the random variable corresponding to total demand satisfied is normally distributed, since it is the convolution of normally distributed market demands. However, the profit function $G(Q^*, y^*)$ is not normally distributed. We ran several tests using simulation to estimate the shape of this distribution. Figure 1 presents the results of 3000 profit realizations of $G(Q^*, y^*)$ in order to approximate the profit distribution. Since there are penalties for underages (e) as well as overages (v), extremely low or high demand

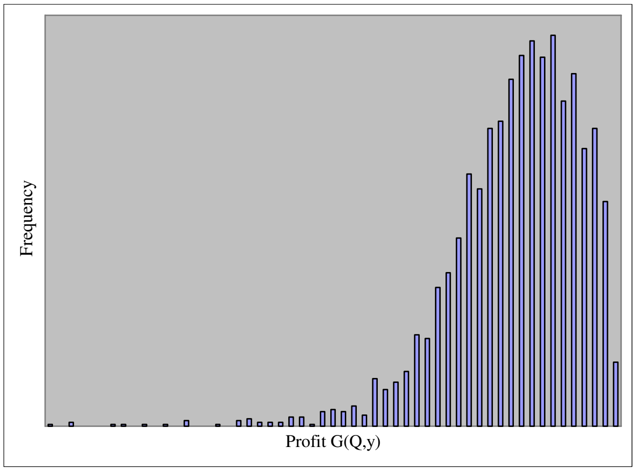

Figure 1: Distribution of Profit Based on Q and y

realizations will result in lower profit (or possibly a loss). These extreme conditions contribute to the left tail of the profit distribution in Figure 1. Notice that the maximum achievable profit does not greatly exceed the expected profit (i.e., it does not have a similar right tail on the distribution). In the newsvendor problem, the critical fractile defines the point in the demand distribution for D^y at which we maximize expected profit. As realized demand moves away from the demand quantity associated with this point, the firm's profit will decrease. However, in our selective newsvendor framework, we also have market-specific revenues r_i associated with each market i. Thus, the maximum profit that the firm can achieve occurs when all realized demand occurs in the market(s) with the highest revenue, and total realized demand still equals the order quantity. Thus, the maximum profit shown in the profit distribution is more well-defined than the maximum loss.

Now consider that the firm would like to minimize the worst-case set of profits (losses). Since the profit distribution is not normally distributed, this complicates the solution approach. In the next section, we show how we can still utilize the fact that the demands are normal in solving the risk averse selective newsvendor problem. Then, in the following section, we consider a more general problem in which demands are not necessarily normal, and we offer a solution approach using simulation to approximate and evaluate potential profit distributions. Bertsimas and Thiele (2003) also consider a data-driven approach, whereby they build upon the sample of available data instead of estimating the probability distributions. They also use individual profit realizations in determining a risk policy to implement, however they only consider one demand distribution.

3 SELECTIVE NEWSVENDOR MODELS WITH RISK

When a firm is concerned about the risk of potential losses, there are many ways in which the firm can actually quantify this risk into a model. We will suggest two classifications of risk models from which a firm may choose.

Our first risk model is based on the original approach defined for the SNP. In the SNP approach, we maximized expected profit, which was the difference between expected net revenue and the cost of demand uncertainty. Now, we consider our firm to be "risk averse" if they still focus on maximizing expected profit, but now they require that there be no more than a certain percentage of potential profits (or losses) below a pre-defined value. We will refer to this value as a defined profit level throughout the remainder of the paper, although a negative value would obviously represent a loss.

We now present the risk averse selective newsvendor as

[RA]

$$\begin{aligned} \text{maximize} \quad & \hat{G}(Q, y) \\ \text{subject to:} \quad & F_G(P) \leq \alpha, \\ & Q \geq 0, \\ & y_i \in \{0, 1\} \quad i = 1, \ldots, n, \end{aligned} \tag{2}$$

where P represents the critical profit value, $\alpha * 100\%$ denotes the percentage of allowable realizations below P, and F_G denotes the distribution of the profit equation $G(Q, y)$. Of course, if we set $\alpha = 0$, this implies that we will not accept any potential for profit less than an amount P. In other words, P represents the minimum acceptable profit level.

Alternatively, the firm may be more focused on risk than on profit. In this case, assume the firm is "risk minimizing," whereby the firm minimizes the worst-case realizations of profit based on a given set of demand sources. We present the risk minimizing selective newsvendor as

[RM]

$$\begin{aligned} \text{minimize} \quad & F_G(P) \\ \text{subject to:} \quad & Q \geq 0, \\ & y_i \in \{0, 1\} \quad i = 1, \ldots, n. \end{aligned}$$

Again, P represents a threshold profit value. Recall that, by adding markets, we can increase our expected revenue (and, most likely, profit), but not necessarily reduce the overall risk. While this may be desirable using model [RA], it is not desirable under model [RM]. The threshold profit level P now becomes the critical factor in determining the preferred market selection set. Also note that the firm must set P such that market will actually be selected. Consider that, for $P < 0$ and $y_i = 0$ for $i = 1, \ldots, n$, we have $F_G(P) = 0$, and we would have an optimal solution with no markets selected. By selecting a value of $P > 0$, however, $F_G(P) = 1$ when no markets are selected, so the model would attempt to add markets to lower this percentage.

3.1 The Critical Profit Level

In this section, we will approximate the value of $F_G(P)$ using simulation. Given a market selection vector $\hat{y}$, an associated order quantity $\hat{Q}$, and a pre-defined critical profit level P, we can estimate the distribution of $F_G(P)$, or the percent of critical profit realizations below P. As stated previously, Figure 1 presents the form of the distribution. Here, we now specify the critical P, and by simulating demand realizations, we can then determine how many of these realizations (or occurrences) will result in a profit below P.

In order to evaluate either model [RM] or [RA], we require this $F_G(P)$ value for every market selection and order quantity tested. For every call to simulation, there will be an associated expense in computational time. Thus, we must be aware of this and try to ration the number of replications performed that still provide an adequate answer in a reasonable amount of simulation time.

3.2 The Optimal Order Quantity

As previously stated, for the selective newsvendor problem, $Q_y^* = F_y^{-1}(\frac{e-c}{e-v})$, or what we can call Q_1. However, for models involving risk, it is not clear that this value should be used for ordering. Figure 2 presents the relationship between the value of Q and the percent of observations not meeting some critical profit level P (i.e., probability that realized profit does not meet some threshold profit level).

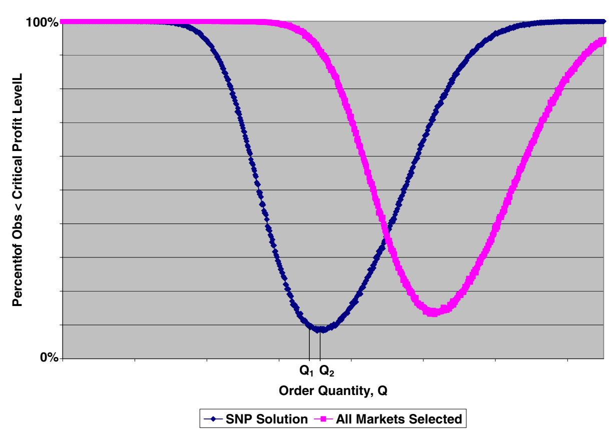

Figure 2: Order Quantity vs. Profit Realized

Note that the value of Q_1 given by the selective newsvendor does not coincide with the order quantity Q_2 that provides the highest probability of meeting our value for P.

Based on the figure, we can also observe that the function describing Q and $F_G(P)$ certainly appears to be unimodal. We will use this conjecture to find the preferred order quantity to use for a given market selection vector $\hat{y}$. If the relationship is indeed unimodal, we can implement a line search technique to converge on the best value for Q. We have chosen to use the golden section technique (see Bazaraa, Sherali, and Shetty (1993)) for our approach. As with the number of replications performed in a simulation, the stopping criterion for finding an appropriate Q will have a direct effect on the overall solution time. If small improvements in $F_G(P)$ require another iteration (and subsequent update in the value for Q), the required number of iterations for convergence will, of course, increase. At each iteration in the line search process, we are performing a full set of simulation replications, which will quickly drive a longer overall solution time.

3.3 Solution Approach

We are now prepared to present solution algorithms for each model, [RM] and [RA]. First, we note the following characteristic concerning the profit contribution of each market to the overall profit. Given some set of selected markets $\hat{y}$, if the addition of market i into the solution reduces the frequency of profits below the critical profit level P (or $F_G(P)$), we would expect this market to be beneficial. We desire such shifts in the profit distribution that reduce the location and size of the left tail of the profit distribution (refer back to Figure 1). We use this idea in constructing the algorithms.

For every new solution (or selection of markets) evaluated, we must perform two main tasks: 1) conduct multiple replications of profit replications via simulation to appropriately represent the distribution of profit; and 2) implement a line search technique to locate a preferred order quantity. In order to appropriately represent profit realizations according to $G(Q, y)$, we will use simulation analysis to populate a profit distribution . With this in mind, we propose the following heuristic procedures to find approximate or near optimal solutions to the SNP models with risk that have been discussed in Section 3. These generic procedures are actually independent of the underlying demand distributions under consideration, although we will focus on markets in which the demand data is normally distributed.

3.3.1 Solving Problem [RM]

We begin with the solution approach to problem [RM], the risk minimizing selective newsvendor formulation. First, we evaluate $F_G(P)$ for every potential market i when i is the only selected market. That is, for every i, we set $y_i = 1$ and $y_j = 0$ for all $j \neq i$, and determine the value of the distribution function of profit, denoted as $F_{G(Q,i)}(P)$. We then re-index all markets $i = 1, \ldots, n$ in non-decreasing order of the value $F_{G(Q,i)}(P)$. Then, starting with re-ordered market [1], we systematically add each market to the solution (i.e., Q, $\hat{y}$), testing for each iteration whether the value of $F_{G(Q,\hat{y})}(P)$ decreases further. The final solution will contain the markets for which a minimum value of $F_{G(Q,\hat{y})}(P)$ is achieved. As with the order quantity graph shown in Figure 2, the selection of markets also behaves similarly. As we add markets, the value of $F_{G(Q,\hat{y})}(P)$ decreases, achieves a minimum, and then begins to increase with additional market selections. This is only provided as an observation. We are presently investigating what can be proven with regards to this relationship. We present the heuristic solution approach to problem [RM]:

Heuristic Solution to [RM]

0. Set $j = 1$.
1. First we select only market j and find the optimal order quantity Q^j for this market selection. We find Q^j based on the line search method proposed in Section 3.2. During the procedure for finding Q^j, we also populate the profit distribution associated with solution vector (Q^j, y^j) using simulation. Then calculate the percentage of worst-case profits for this market selection, previously defined as $F_{G(Q,j)}(P)$.
2. Update $j = j + 1$; Repeat Step 1 until $j > n$.
3. Sort the markets in non-decreasing order of $F_{G(Q,j)}(P)$ values to obtain the sorted market order [1],[2],[3],...,[n]. Reset j to $j = 1$. We will call this sorted market order the preference order for selecting markets in order to minimize $F_G(P)$.
4. Select markets [1],[2],...,[j] and estimate $F_G(P)$ for this market selection, denoting this value as $F_G(P)[j]$. Again, in order to find $F_G(P)[j]$, we must populate the profit distribution using simulation. Similarly, we find $Q[j]$ using the line search method proposed in Section 3.2.
5. Update $j = j + 1$; Repeat Step 4 until $j > n$.
6. We calculate n such potential solutions to problem [RM]. From the set $S = \{F_G(P)[j], j = 1, \ldots, n\}$, the solution to [RM] would be the least value in this set.

Note that our solution approach does not require that we investigate all 2^n possible market selections, which would be computationally prohibitive (as we will show in our computational tests in Section 4. Instead, we only search n potential solutions.

We now propose a similar heuristic algorithm for solving the risk averse selective newsvendor problem, or problem [RA].

3.3.2 Solving Problem [RA]

In the case of problem [RA], we still proceed with a similar solution approach with the exception of how the algorithm begins. We first compute the optimal solution to the selective newsvendor problem (SNP) with no risk aversion using the Decreasing Expected Net Revenue to Uncertainty (DERU) Property, first presented in Taaffe et al. (2004). If the SNP solution satisfies risk averse constraint (2), we conclude that the SNP solution is also the optimal solution to problem [RA]. Otherwise, we check other potential market selections, and the associated new order quantity, against the risk averse constraint. When the constraint is satisfied, we test to see if the expected profit generated exceeds the incumbent solution value, and the incumbent is updated as necessary. Finally, based on our risk averse constraint (2), if $F_G(P) > \alpha$ for a particular solution, this solution can be immediately eliminated from consideration. We present the heuristic solution approach to problem [RA]:

Heuristic Solution to [RA]

1. Sort the markets in non-decreasing order of the DERU Ratio Property (see Taaffe et al. (2004)) to obtain the sorted market order [1],[2],[3],...,[n]. Denote $\hat{G}(Q, y)$ as the expected value of the profit equation described in equation (1). Obtain the optimal solution to the SNP problem without risk, which will contain all markets in the sorted list up to and including some market $j = [1], [2], \ldots, [n]$. Call this solution $\hat{G}(Q, y)[0]$. Update $Q[0]$ based on the line search method proposed in Section 3.2. Populate the profit distribution such that the percentage of worst-case profits $F_{G(Q,y)}(P)[0]$ can be estimated. If constraint (2) is satisfied, STOP with an optimal solution equal to the SNP solution. Otherwise, set $j = 1$ and continue.
2. Select markets [1],[2],...,[j]. Estimate $F_G(P)$ for this market selection, denoting this value as $F_G(P)[j]$. Also calculate the expected profit of this solution as $\hat{G}(Q, y)[j]$. Appropriate values of $F_G(P)[j]$ and $Q[j]$ are found via simulation and line search as previously discussed.
3. Update $j = j + 1$; Repeat Step 2 until $j > n$.
4. We calculate n such potential solutions to problem [RA]. From the set $S = \{F_G(P)[j], \hat{G}(Q, y)[j], j = 1, \ldots, n\}$ the solution for the [RA] model would be the greatest value of $\hat{G}(Q, y)[j]$ for which its corresponding value of $F_G(P)[j]$ is feasible.

Again, note that our solution approach does not require that we investigate all 2^n possible market selections. Based on a limited set of problem instances, we see that the heuristic actually achieves the SNP solution without risk in all cases. This is primarily due to the following observation. The difference between the minimum α found in [RM] and the α determined based on an SNP solution approach are quite similar. This leads to the solutions to [RM] and [RA] being the same.

4 COMPUTATIONAL TESTS

We present two sets of results that describe the performance of the algorithm, as well as the change in solution values from the original SNP solution. For both models, we will assume that the pre-defined profit level P can be calculated as 10% of the expected profit given by the SNP solution approach.

We created a small set of 20 test problems or instances for each size of the potential market pool: 5, 10, and 20, respectively. Every market has unit revenue in the range U[\$200,\$240], while the unit production cost is set at \$200. Expected demand and demand variance for each market are distributed according to U[500,1000] units and U[50000,100000], respectively. The fixed cost for market entry are drawn from U[\$2500,\$7500]. Finally, the salvage value is \$50 per unit, and the expediting or shortage cost \$500 per unit, respectively. For each simulation replication or demand realization, we calculate demand based on the market selection variables y_i for that particular solution. We also use these selection variables when calculating expected profit values during the solution of problem [RA].

At 1000 and 5000 simulation replications, we still experienced significant variability in the value of $F_G(P)$. Thus, we opted to run 10,000 replications for each solution tested. The golden section line search technique evaluated order quantities in a range of 0 to the maximum total demand if all markets were included (and demand in each market was realized at its highest level). The procedure would converge on an order quantity once the current best quantity produced less than a 1% improvement from the prior iteration's order quantity value.

In conducting the experiments, we quickly realized that solving each problem using full enumeration was extremely costly. Combining the requirement of simulation and line search, the solution time for a 20-market problem easily exceeded several days. We had the choice of either reducing the number of replications for our profit function, or we could reduce or eliminate the line search. Since we observed significant variations in $F_G(P)$ values at smaller replication counts, we chose to remove the line search technique to reduce solution time. In its place, we used the preferred order quantity generated via the SNP approach. Thus, the order quantity used (when line search was not performed) was $Q^y = \sum_{i=1}^n \mu_i y_i + z(\rho)\sqrt{\sum_{i=1}^n \sigma_i^2 y_i}$. Table 1 presents a comparison of the full enumeration approach and our proposed heuristic approach to solving problem [RM]. In each approach, we implemented the respective algorithms with line search (LS) and without line search (NLS) to determine a preferred order quantity.

Table 1: Comparison of Solution Approaches to [RM]

Scenario/ Measurement	Solution Approach: Enumeration LS	Enumeration NLS	Heuristic LS	Heuristic NLS
5 Markets				
CPU Time	11 sec	<1 sec	5 sec	<1 sec
$F_G(P)$	0.281	0.288	0.284	0.288
Q	3052	3153	3115	3051
10 Markets				
CPU Time	12 min	29 sec	17 sec	<1 sec
$F_G(P)$	0.1281	0.1284	0.127	0.128
Q	5837	5953	5829	5837
20 Markets				
CPU Time	NA	14.9 hr	1 min	2 sec
$F_G(P)$	NA	0.1481	0.1529	0.1519
Q	NA	12091	12145	11650

Notice how similar the results are between the two solution methods. In fact, our heuristic method actually outperforms the enumerative approach in some cases. (Recall that the enumerative procedure is still a heuristic itself, since we must use simulation to construct the profit distribution for every potential market selection assignment.) Thus, we can at least say that we are not giving up much in the way of solution quality for a significant reduction in solution time.

Line search, in general, appears to improve solution quality, but this is not consistent across all test cases. We observe solution improvements for all 5- and 10-market cases, but not for the 20-market case, where the minimum percentage of worst-case profits (when using our heuristic) is found without line search. A caveat on the 20-market case is that it only contains two test instances due to the long solution times using the enumerative approach.

There does not appear to be one overall relationship between the value of Q and the choice of line search or no line search. In some instances, the order quantity found via line search was smaller than its counterpart, and in other cases, it was larger. Additional insight into a possible relationship will require more analysis.

We also note that, for cases when the same market selection vector was selected (which did occur quite often), the objective value was actually different. This can be contributed to the randomness that exists in the simulation process to calculate the profit distribution, despite performing 10,000 replications. If each problem does not contain identically-shaped profit distributions, then it follows that the minimum percentage of worst-case profits will not be the same even when the market selection is the same.

Next, we present the results for a similar solution comparison for problem [RA]. In problem [RA], the firm is trying to maximize expected profit while meeting a risk criterion that restricts the allowable percentage of profit observations below some critical level P. In order to create this constraint (2), first let α_{SNP} represent the $F_G(P)$ value found with the SNP solution without risk. Then the α used in the risk constraint is $\alpha = 0.95 * \alpha_{SNP}$. Table 2 presents the results of the comparison for problem [RA]. The additional statistic, $\hat{G}$, represents the expected profit objective of problem [RA].

Table 2: Comparison of Solution Approaches to [RA]

Scenario/ Measurement	Solution Approach			
	Enumeration		Heuristic	
	LS	NLS	LS	NLS
5 Markets				
CPU Time	11 sec	<1 sec	3 sec	<1 sec
$\hat{G}$	36296	39895	39895	34286
$F_G(P)$	0.1514	0.1552	0.1546	0.1409
Q	3125	4514	4476	2314
10 Markets				
CPU Time	17 min	26 sec	8 sec	<1 sec
$\hat{G}$	66627	67073	66839	65581
$F_G(P)$	0.1092	0.1048	0.1071	0.1083
Q	6714	6926	7041	6582
20 Markets				
CPU Time	NA	13.5 hr	26 sec	2 sec
$\hat{G}$	NA	80963	80963	80963
$F_G(P)$	NA	0.1617	0.1539	0.1637
Q	NA	12419	12779	12419

For problem instances in which the solution to $\hat{G}_{SNP}$ also represents one with minimum profit risk ($F_G(P)$), using α instead of α_{SNP} in the constraint will make the problem instance infeasible. This occurred several times for each market class. We only report results on those instances in which we could find a feasible solution. As with problem [RM], we see the similarity in objective values ($\hat{G}$, in this case) across each solution approach. However, there is more variability in this case. The expected profit values are for these problems are consistently below those generated for $\hat{G}_{SNP}$, since we force the model to find a solution with less risk (i.e., $\alpha = 0.95 * \alpha_{SNP}$).

We found similar results as for problem [RM] concerning the relationship between Q, $F_G(P)$, and the use of line search or no line search. Overall, we achieve similar quality solutions using our heuristic approaches as compared to the enumerative approaches, with the noise in solution quality due to the simulations required to develop the profit distribution. We will be examining this relationship in more detail in future work.

For the 20-market case without line search, we observe that the same solutions were found for both the heuristic and enumeration approaches, with an average order quantity of 12419, but the $F_G(P)$ values are slightly different. As in the discussion of Table 1 results, any difference in the shape of the profit distribution will cause slight variations in this fractile value. The order quantities and expected profits, however, are identical since neither depends on the variation in demand values.

5 CONCLUSIONS

After completing our experiments, we were able to conclude that our proposed heuristics for solving the selective newsvendor models with risk offer high quality solutions at a fraction of the time of an enumerative approach. Based on initial testing, the order quantity was neither sensitive to the objectives of the risk models nor the line search implementation. If we were to consider higher critical profit levels, in which more scenarios or outcomes would fall below the critical value, altering the value of Q may have more effect, but additional testing would need to be performed to truly understand this relationship.

We point out that obtaining solutions to probabilistic risk models can be quite cumbersome, and we offer approaches that firms dealing with risk issues can implement. We are presently evaluating a theoretical approach to the risk problem when demand data are normal. However, when we know our demands are not normal, such as the case when market demands will either be realized in their entirety or not at all (i.e., the so-called all-or-nothing demand problem), we must still resort to an approach using simulation as described in this research. We plan to investigate the relationship between various types of demands and the applicability of the procedures outlined in this work.

REFERENCES

Bazaraa, M., H. Sherali, and C. Shetty. 1993. *Nonlinear programming: Theory and algorithms*. 2nd ed. New York, New York: John Wiley & Sons.

Bertsimas, D., and A. Thiele. 2003. A robust optimization approach to supply chain management. Working paper, Massachusetts Institute of Technology.

Cachon, C. 1999. Competitive supply chain inventory management. In *Quantitative Models for Supply Chain Management*, ed. S. Tayur, M. Magazine, and R. Ganeshan. Boston, Massachusetts: Kluwer Academic Publishers.

Carr, S., and W. Lovejoy. 2000. The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. *Management Science* 46 (7): 912–927.

Chen, X., M. Sim, D. Simchi-Levi, and P. Sun. 2004. Risk aversion in inventory management. Working paper, Department of Mechanical and Industrial Engineering.

Collins, R. 2004. The behavior of the risk-averse newsvendor for uniform, truncated normal, negative binomial and gamma distributions of demand. Working paper, Santa Clara University, Department of Operations and Management Information Systems.

Eeckhoudt, L., C. Gollier, and H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Science* 41 (5): 786–794.

Lau, H. 1980. The newsboy problem under alternative optimization objectives. *Journal of the Operational Research Society* 31:525–535.

Petruzzi, N., and M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations Research* 74:183–194.

Petruzzi, N., and G. Monahan. 2003. Managing fashion goods inventories: Dynamic recourse for retailers with outlet stores. *IIE Transactions* 35 (11): 1033–1047.

Porteus, E. 1990. Stochastic inventory theory. In *Handbooks in Operations Research and Management Science; Volume 2*, ed. D. Heyman and M. Sobel, Chapter 12, 605–652. Amsterdam, The Netherlands: North-Holland.

Schweitzer, M., and G. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution. *Management Science* 46 (3): 404–420.

Taaffe, K., J. Geunes, and H. Romeijn. 2004. Market entrance and product ordering decisions under demand uncertainty. In *Proceedings of IIE Research Conference*. Houston, Texas.

Taaffe, K., J. Geunes, and H. Romeijn. 2005. Integrated marketing and operational decisions: The selective newsvendor problem. Working paper, Department of Industrial Engineering, Clemson University.

Taaffe, K., and H. Romeijn. 2005. Selective newsvendor problems with all-or-nothing order requests. In *Proceedings of IIE Research Conference*. Atlanta, Georgia.

Tsay, A., S. Nahmias, and N. Agrawal. 1999. Modeling supply chain contracts: A review. In *Quantitative Models for Supply Chain Management*, ed. S. Tayur, M. Magazine, and R. Ganeshan. Boston, Massachusetts: Kluwer Academic Publishers.

Van Mieghem, J. 2004. Risk-averse newsvendor networks: Resource sharing, substitution, and operational hedging. Working paper, Northwestern University, Kellogg School of Management.

AUTHOR BIOGRAPHIES

KEVIN TAAFFE is an assistant professor in the Department of Industrial Engineering at Clemson University. His primary research interests include production, inventory, and demand management, as well as transportation and logistics systems analysis. He is a Senior Member of IIE, as well as a member of INFORMS, Transportation Research Board, and Decision Sciences Institute. His email address is <taaffe@clemson.edu> and his Web address is <http://people.clemson.edu/~taaffe/>.

DEEPAK TIRUMALASETTY is a graduate student pursuing a Master of Science degree in the Department of Industrial Engineering at Clemson University. He has research interests in both production and inventory systems, as well as transportation systems. His email address is <dtiruma@clemson.edu>.