

USING SIMULATION TO EVALUATE A CONTINUOUS REVIEW (R, Q) TWO-ECHELON INVENTORY MODEL

Yeu-San Tee and Manuel D. Rossetti, PhD¹,

Department of Industrial Engineering
University of Arkansas
4207 Bell Engineering Center
Fayetteville, AR 72701

Abstract: The effective management of supply chains depends upon the ability of decision makers to control inventory while maintaining customer service at low cost. This paper examines an analytical model for the two-echelon one warehouse and multiple retailers system using (R, Q) inventory policies. A simulation model was developed to analyze the performance of the analytical model under conditions that violate the model's fundamental modeling assumptions. The behavior of the multi-echelon inventory model under different demand (non-Poisson) and lead-time (non-constant) conditions were observed in an experimental design. The simulation results showed that there was significant error in estimating the total system cost (tends to be underestimated) when the model assumptions are violated. The cost relative error ranged between 34% and – 40% over all the experiments examined. Models of the one warehouse and multiple retailers two-echelon inventory system are significant because they can be embedded within aggregation/disaggregation techniques that allow the general approximation of supply chain networks. Practitioners looking to utilize these models will have a better understanding of when and when not to use these models in practice.

1. INTRODUCTION

Efficient and effective management of inventory throughout the supply chain can significantly improve customer service levels and reduce cost. Many multi-echelon inventory models are used to determine the decision rules (inventory policies) for jointly controlling the inventory levels at all locations in a distribution system. These models are sometimes integrated into enterprise resource planning systems and are treated as black boxes when obtaining solutions. These models should be robust enough to provide accurate and useful results over a wide variety of conditions found in practice. This research examines the robustness of a (R, Q) continuous review two-level inventory model via simulation.

The distribution network under consideration consists of one warehouse and N retailers, where the retailers directly serve the customers and the warehouse supplies all the retailers. At each facility location, when the inventory position (net inventory on hand plus stock on order minus backorders) drops below the reorder point R , a replenishment order quantity of Q is placed, i.e. the (R, Q) continuous review replenishment policy. Many have suggested using the (R, Q) inventory control policy on the slow moving type A items or in situations where the ordering cost is high [Silver, Pyke, and Peterson (1998), Hopp and Spearman (2000), Zipkin (2000), Axsater (2000), etc.].

Axsater (2000) presents a procedure for evaluating two-echelon inventory systems under stochastic conditions. The work is significant in that exact analytical results are obtained for the model and the model can then be embedded within aggregation/disaggregation procedures for evaluating supply chains. The model is built under certain assumptions. This study investigates the effects of violating those modeling assumptions, and determines the conditions where the model will perform poorly. The total system cost is an important measure for this study, and it serves as the basis for comparisons. Only if the model is proved to be robust should one suggest its general use in practice. In addition, the understanding gained will help develop more robust models in the future. To examine the robustness of the model, simulation experiments over a wide variety of parameters were performed to estimate the true system cost and allow comparison to the cost predicted by the model. The simulation model built in this study is simple and ease to use, and the model can be modified to accommodate other distribution systems.

The next section provides an extensive review of relevant literature concerning multi-echelon inventory systems and robustness studies. The methodology is explained in the following sections where we present the reconstruction of the analytical model, the building of the simulation model, the design of experiments, and the results. The last section concludes with recommendations on directions for future research.

¹ Corresponding Author (Email: rossetti@uark.edu, Fax: 501-575-8431)

2. LITERATURE REVIEW

2.1 Multi-echelon Inventory Control

Many early multi-echelon inventory models have been used for contingency support, such as the work of Sherbrooke (1968) and Muckstadt (1973) on controlling the reparable items in the military base-depot supply system. Since then, many other multi-echelon inventory systems have been studied extensively, especially for service part inventory control and for the one-for-one base stock ordering policy (a special case of (R, Q) policy with $Q = 1$). Previous research has led to the development of analytical inventory models. Of particular note is the work of Cohen et al. (1986), which has been successfully integrated into IBM's OPTIMIZER, a multi-echelon service inventory optimization software support system. IBM reported a savings of over \$250 million resulting from the use of their OPTIMIZER software.

The continuous review (R, Q) policy two-echelon system has also received tremendous attention from researchers. A good review of the models dealing with continuous review policies for multi-echelon inventory systems can be found in Axsater (1993a). The optimal solution of a two-echelon inventory system depends on the assumptions used in the model: system structure, cost function, demand distribution, replenishment lead-time, etc. The research in this area can be categorized into two main approaches. The traditional method focuses on the steady-state behavior of the inventory levels, where the lead-time demand is approximated by the mean and variance of a certain distribution. The multi-level system is decomposed into single locations to be evaluated separately with parameters that depend on each other. The total cost function is obtained through the average inventory and backorder units. A seminal paper using this approach is Deuermeyer and Schwarz (1981). Deuermeyer and Schwarz's work was followed by Moinzadeh and Lee (1986), Lee and Moinzadeh (1987), and Svoronos and Zipkin (1988). The approximation method of Svoronos and Zipkin (1988) has been shown by Axsater (1993a, 1993b) to be very accurate for the identical retailer case.

As opposed to the traditional approach, Axsater (1993a, 1993b) suggested several approximation methods. These approximations were derived based on the previous work of Axsater (1990), which is a recursive procedure for the evaluation of one-warehouse and N retailers utilizing one-for-one policies. The holding and backordering costs must be functions of the delay experienced by the customer. His numerical results show that his approximations provide good results that are comparable to that of Svoronos and Zipkin (1988). Along the line of expressing the total system costs as a weighted average of one-for-one policies, Axsater (1998) extended his previous model to evaluate the non-identical retailer case. He successfully developed the exact solution for the case of two retailers, and then used that solution to approximate for the case of more than two retailers.

Forsberg (1996) redefined the probability distribution functions that are used to determine the weights for the weighted mean costs for one-for-one policies. As a result of this new cost expression, the two-level inventory system with one-warehouse and N non-identical retailers can be evaluated exactly under their assumptions. Forsberg (1997) also extended this exact evaluation to the case in which the demand process is not Poisson distributed. His model provides good approximations for the system with Erlang distributed customer inter-arrival times.

Since previous models by Axsater and Forsberg were based on the weighted average costs for one-for-one policies, the models are limited to pure Poisson demand processes only and rely on a cost structure that is not commonly used. Axsater then investigated the steady-state behavior of the inventory levels. He first developed an approximation method to evaluate the two-level system with non-identical retailers under the compound Poisson demand process (Axsater, 1995). The cost structure in this approximation was still constrained to be a function of delay and storage time units; however, it did not depend on the costs of one-for-one policies.

Finally, Axsater (2000) provided an exact analysis through determining the complete probability distributions of the retailer inventory levels in steady state. This model uses a common cost structure that can more easily be understood, and the model can be used to solve the one-warehouse and non-identical retailer system with compound Poisson demand. This model, by far, is the state-of-art in exactly evaluating two-echelon inventory systems with continuous review (R, Q) batch ordering policies for the low demand items such as spare parts. Therefore, the Axsater (2000) was selected for testing in this study.

2.2 The Robustness of Inventory Models

Previous studies have been performed to test the robustness of inventory models directly and indirectly. Those studies were done using several formats and approaches. Some used analytical models as opposed to computer simulation. Some

evaluated the system performance through a sensitivity analysis on several parameters, and some studied the influence of just a single parameter in the model, especially the demand factor. Many of the studies were performed on single location

inventory models. The results of Naddor (1978), Fortuin (1980), Banks and Spoerer (1986), and Lau and Zaki (1982) conclude that solutions are affected more by the means and standard deviations of the demands rather than the form of the demand distribution. The effect of the demand and lead-time distribution was found significant only when the differences between the variances (standard deviations) of the probability distributions are large.

The normal approximation has been often used as an approximation for the lead-time demand distribution, especially when determining the safety stock. Tyworth and O'Neill (1997) examined the use of such approximations in (R, Q) continuous review inventory models. By comparing the solutions from the normal theory and exact approaches, they found that the normal approximation method is robust across seven industry groups (fast-moving demand items). Fotopoulos, Wang, and Rao (1988) presented a method to determine the safety stock when the demands are autocorrelated and the lead-times are random. Their numerical results show that ignorance of autocorrelation in demand could provide severe errors; however, the effect of non-normal demand is found to be relatively small.

The above studies were all concerned with single-echelon inventory models. Lagodimos, De Kok, and Verrijdt (1995) performed a study on multi-echelon models. They tested the robustness of two two-echelon serial inventory models under stationary autocorrelated demand processes. The essential differences between their research and this one are that both of the models examined in their research assumed independent identically distributed normal demands, and the models operate with a periodic review order-up-to-S inventory policy. The approach of their study is analytic and their parameters are redefined to extend the tested models to fit any stationary autocorrelated normal demand process in an exact closed form. Their results demonstrate that the models, which ignore the effects of autocorrelation, might provide significant error in the system performance measures depending on the overall system parameters setting.

Our review of the literature has indicated that no studies have been performed to examine the robustness of two-level multi-echelon inventory models, especially those of the form given in Axsater (2000). In this research, we are interested in understanding the circumstances under which violations of the model's assumptions cause errors and deteriorate the accuracy of the models. If the deterioration is not significant, then we can be more assured of the model's applicability in practice.

3. METHODOLOGY

Our basic methodology is to compare the results of the simulation under various conditions tested to the results of the analytical model. The optimum (R, Q) values recommend by the analytic models are examined to determine the sensitivity and robustness of the models. First, the solutions and their associated performance measures are obtained by using the analytic models. The solutions then serve as the input for the simulation model. By manipulating the capability of the simulation software, the simulation model is run with various conditions that are different from the assumptions of the analytic model. The simulation model captures the performance measures of the system for comparison.

3.1 The Axsater (2000) Exact Analysis Program

A prototype program is available at <ftp://ftp.ie.lth.se/ie/sven/ExactAnalysis/> for the Axsater (2000) model. The source code of the program was obtained from the author. The program is developed within a Microsoft Excel spreadsheet environment with Visual Basic for Applications, and the application program is compiled in Visual Basic. In order to run the program with larger problems such as more than 10 retailers, the source code was modified to handle larger array sizes. The program was used to solve all 32 test problems that have been suggested by Svoronos and Zipkin (1988), and then the solutions were used as the base cases for testing. The recommended inventory policies were used as the simulation inputs. In these 32 test problems, the system structure and assumptions are:

- The system has one warehouse and N identical retailers each with continuous review (R, Q) inventory policies.
- The demand is assumed to be stationary Poisson process at the retailer level.
- The replenishment lead-time is deterministic.
- All unsatisfied demands are backordered.

3.2 The Simulation Model

In this study, discrete-event simulation models were built to test the robustness of the two-echelon inventory models found in Axsater (2000) and Svoronos and Zipkin (1988). For the conditions involved within this study, we are interested in the steady state performance measures: the expected inventory on-hand at the retailers and at the warehouse, the expected

number of item backordered at the retailers, and the expected total system cost. The research used Arena 5.0 Professional Edition to develop the models. Arena 5.0 is simulation software that utilizes SIMAN simulation language as the main mechanism behind its high-level user-friendly interface (Kelton et. al., 1998).

3.2.1 Model Structure

Since both the warehouse and retailers follow the same type of inventory policy, a single location inventory system was first simulated. The flowchart in Figure 1 illustrates the main inventory control activities at a single facility location. When a demand (customer order) occurs, the unit demanded is determined, and then the system checks for the availability of stock. If the stock on-hand is enough for the order, the demand is filled and the quantity on-hand is decreased. On the other hand, if the stock on-hand is not enough to fill the order, the entire order is backordered. The backorders are accumulated in a queue and they will be filled on a first-come-first-serve basis after the arrival of replenishment order. The inventory position (inventory on-hand + on-order – backorders) is checked each time after a regular customer demand and the occurrence of a backorder. When the inventory position falls under the reorder point, a replenishment order is placed. The replenishment order will take a certain time to arrive and fill any backorders at the retailer and increase the on-hand inventory.

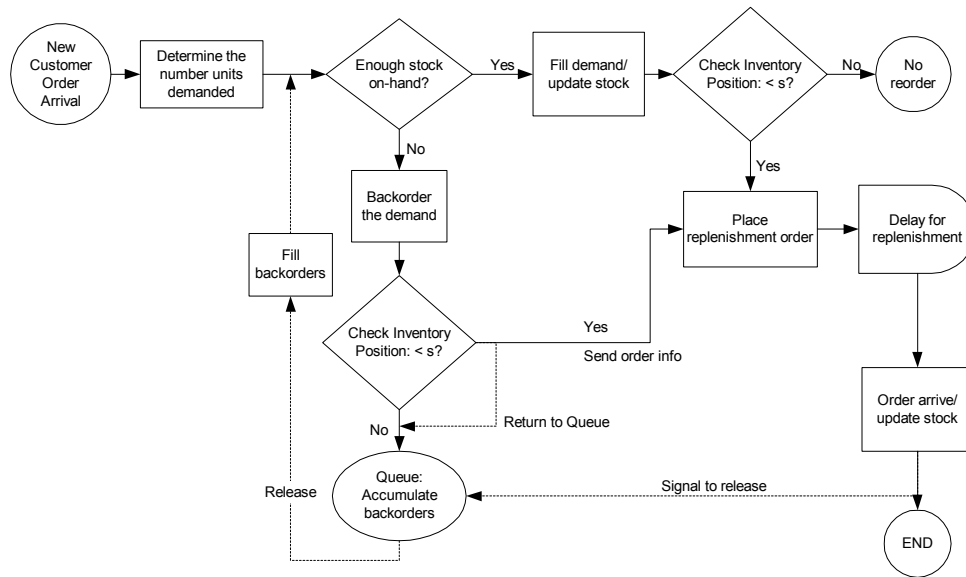


Figure 1. Flowchart of the Single Location Inventory Control Activities.

The above logic in the single location inventory model was expanded into a two-level inventory system (one warehouse and N retailers) by adding in the relationships and interactions between levels. The main control activities at the retailer are the same as the single location model above until the retailer places a replenishment order. The order will be sent directly to the warehouse model. The demand process at the warehouse depends on the order frequency and order quantity at the retailers. The inventory control activities at the warehouse are similar to the retailer. When the demand is filled at the warehouse, the stock is transferred to the retailer. If the warehouse is out of stock, the demand will be backordered, and the retailer will have to wait for extra time to get its replenishment. Figure 2 shows the flow of the orders between the warehouse and retailers in the simulation. Within the simulation model, the N retailers were generalized such that only one set of retailer logic blocks was necessary because of the use of arrays.




Figure 2. The Order Flow in the Multi-echelon Inventory System.

The following assumptions were also incorporated into the model:

1. The demand process at the retailer is established through the specification of the time between arrivals and the demand quantity. A general renewal process with the time between orders governed by a specific probability distribution as given in the experimental design was used. Each demand was assumed to be a single unit of product.
2. All unsatisfied demand will be backordered and no partial filling of an order is allowed.
3. The replenishment lead-time of the warehouse is constant, and supply is unlimited.
4. The retailer lead-time can be deterministic or stochastic. After performing some analysis on the simulation models with stochastic lead-times (gamma and exponentially distributed), the probability that more than a single order is outstanding is very small in our test cases. Therefore, in this study we will ignore the effect of order crossing such that the replenishment orders may cross when multiple orders are outstanding. For further information on how to handle this problem, we refer to Zipkin (1986).

3.2.2 Model Verification and Validation

After the simulation model was built, step-by-step debugging of the Arena simulation program was performed to see if the programming logic matched the conceptual model (see the flowcharts in Figure 1 and 2 above). The simulation model was verified to run as intended. The model also has to be validated so that the performance measures given are an accurate and promising representation of the actual system. The model was first validated for the single location inventory model case. The one retailer model was compared to an analytic model with similar characteristics in Zipkin (2000) (p.190, example 6.2A, part 2). The comparison showed that the simulation model provides accurate measures for the system performance (deviations with less than or equal to 0.01 at standard errors of 0.3 or less under 30 replications of 10 years run-length).

The logic and entity movement in the two-echelon program were also checked and verified. The Axsater (2000) model was used to determine the exact values for the 32 test problems of Svoronos and Zipkin (1988). The exact evaluation of these 32 test problems was compared to the outputs from the simulation model, and the simulation results showed little deviation from the analytical results. The deviations of the simulated total system cost from the exact value were only 0.07 or less (with standard errors of 0.04 or less under 30 replications of 50 years run-length). Therefore, both the single location and two-echelon inventory models were validated under the condition of deterministic lead-times with Poisson demand process at the retailer. The simulation model accurately predicts the performance of the system.

3.2.3 Steady-State Simulation

According to Needham and Evers (1998), an inventory system of this nature is a non-terminating system and one must design the experiments to evaluate the system under steady-state conditions. The experimental analysis of a simulation model must provide sufficient independent observations to do statistical tests and to obtain statistical significance. Initial inventory on-hand, on-order, and backordered at each location were set to zero for the simulation model, and these conditions could cause initialization bias. Welch's plot procedure as described in Law and Kelton (2000) was used to determine the warm-up period. The warm-up period of five years was used to mitigate the effect of initialization bias. The

batch size of 50 years was found to be sufficient to produce uncorrelated batch means and unbiased variance estimates from the normality and time series autocorrelation tests. Thirty batches were determined to be adequate to ensure a 95% confidence interval width of less than 0.34 for the average annual total cost measurement. As a summary, a warm-up period of five years and 30 batches of 50 years were used for all experiments within the study.

4. EXPERIMENTS AND RESULTS

Factorial experiments are designed to evaluate the effect of each individual factor, and the possible interactions between factors. The behaviors of the multi-echelon inventory model under different demand and lead-time conditions are of interest in this study. In particular, we examined the effect of violating the Poisson arrival process assumption at the retailers, and we examined the effect of non-constant lead-time. For the renewal arrival process, the gamma distribution was chosen due to the fact that the exponential distribution is a special case of the gamma. As for the lead-time distribution, many have found that the gamma distribution is desirable in the real world [Bagchi, Hayya, and Ord (1984) and Tyworth and O'Neill (1997)].

4.1 Experimental Design

The experiments consisted of a total of seven factors as shown in Table 1. The original 32 test problems given in Svoronos and Zipkin (1988) consisted of the first five factors. We added the coefficient of variation ($C_v = \text{standard deviation} / \text{mean}$) of the demand and lead-time factors to control the variability of the demand and the lead-time. Six scenarios were run for each of the original 32 test problems. Therefore, the factorial design consisted of a total of 192 test cases (design points).

Table 1. Experimental Factors and Levels.

Factors	Symbols	Levels	Level
Average Demand	D	2	Low - 0.1 unit demand per period
			High - 1.0 unit demand per period
Number retailers	N	2	Small - 4 retailers
			Large - 32 retailer
Backorder cost factor	P_r	2	Small - \$5 per unit backordered
			Large - \$20 per unit backordered
Retailer order quantity	Q_r	2	1 unit
			4 units
Warehouse order quantity	Q_w	2	1 Q_r batch-unit
			4 Q_r batch-units
Demand C_v	D_{C_v}	2	Gamma distributed (low $C_v = 0.4$)
			Gamma distributed (high $C_v = 1.6$)
Lead-time C_v	L_{C_v}	3	Gamma distributed (low $C_v = 0.4$)
			Gamma distributed ($C_v = 1.0$, exponential)
			Gamma distributed (high $C_v = 1.6$)

Since we are interested in measuring the robustness of the model, we need to select responses that compare the model's performance to that of the simulation. In an inventory system, the main measure of performance is the average total system cost, which is computed from the expected on-hand inventory level at the retailer, the expected number of backorders at the retailer, and the expected on-hand inventory level at the warehouse as defined in the Svoronos and Zipkin (1988). To measure the robustness of these models in estimating the performance of the system, the deviation (error) of the estimated values from the actual system values (simulated values) are needed for comparisons. If the error is a positive number, the model overestimates the system, which means the actual value is lower. If the error is negative, the system performance is underestimated such that the actual value is higher. Are these deviations acceptable to guarantee the use of the model? Under what condition(s) will the model perform better?

Since each test case has different conditions with certain recommended inventory policies, the performance measures are not the same for each case. The error will show how much the model deviates from the actual value. The error and relative error are defined as follows:

$$Error = (Model\ Estimated\ Value) - (Actual\ Simulated\ Value)$$

$$Relative\ Error = Error / (Actual\ Simulated\ Value)$$

Nevertheless, the relative error has a disadvantage, which is the over sensitivity of this measure. When the system has small performance measures, little deviation will give a high relative error. The error and relative error of the average total system cost are the main responses used in the statistical analysis of the factorial experiments in this study. In these experiments, the 192 cases were run with 30 batches yielding 30 estimates of each performance measure. The simulation captured a total of 5760 observations.

4.2 Results and Discussions

Table 2 reports the summary statistics of the error and relative error in predicting performance measures under all testing conditions from all the simulation observations (across all experiments). The summary statistics include the sample mean, sample standard deviation, maximum, upper quartile, median, lower quartile, minimum, and sample size. Note: one should not assume a direct relationship between the error and relative error summary statistics. There is no guarantee that the arrangement of the error and relative error is in the same order. For example, the maximum relative error might not come from the maximum error due to the over sensitivity issue of the relative error mentioned above.

Table 2. Overall Total System Cost Error and Relative Error Summary Statistics.

	Average Total Cost	
	Error	Relative Error
Mean	-4.16	-0.04
Std Dev	12.34	0.13
Maximum	4.53	0.34
Upper quartile	0.43	0.04
Median	-0.60	-0.03
Lower quartile	-3.16	-0.11
Minimum	-75.99	-0.40
Sample Size	5760	5760

In Table 2, the sample means of the error and relative error for the total cost across all experiments are negative values, which indicate that the model has a tendency to underestimate the total cost under the testing conditions. This is further supported by the negative median values in that fifty percent or more of the data points for the total cost predicted are underestimated. The purpose of our experiments was to assess the quality of the model estimates on the performance measures relative to the true value when the assumptions are violated. The cost is overestimated as much as 34% from its true value (maximum relative error), and the cost is underestimated as much as 40% (minimum relative error).

In order to examine the risk that the model performs badly, we examined the probability that the model estimation will be greater than 10% of the true value. Table 3 shows the summary for such probabilities from the experiments by counting the number of times the absolute relative error was greater than 0.1 and then dividing by the sample size of 5760. The table also includes the probability when the demand rate is low ($D = 0.1$) and high ($D = 1.0$) each with 2880 samples to show the significance of the demand factor. As shown in Table 3, the model does not perform well in predicting the cost. Of the experiments 38% have error greater than 10% from the true value. When the demand is high, the model performs poorly with 60% of the experiments having a cost within 10% of the true value. On the other hand, when the demand is low, 15% of the experiments have error of greater than 10% of the true value. The non-Poisson demand and non-constant lead-time conditions clearly have a significant effect on the total cost performance measure.

Table 3. Probability of Error Greater than 10% of the True Value from Model Estimate.

Performance Measures	Probability (absolute relative error > 10%)		
	D = 0.1	D = 1.0	Overall
Cost	0.15	0.60	0.38
Sample Size	2880	2880	5760

The ANOVA tests indicated that most of the main factors (except the Q_w factor) significantly affecting the cost error and relative error. The main effect plots for the cost relative error showed that increasing the level of any factor (except the Q_w) from low to high will increase the relative error: the cost is underestimated more when the D , N , P_r , Q_r , D_{C_v} , and L_{C_v} are high. Besides the main effect plots, the interaction plot of D_{C_v} vs. L_{C_v} in Figure 3 below provides important insights about the behavior of the cost relative error under the violated demand and lead-time conditions. When the demand variability is smaller than the model assumption ($D_{C_v} = 0.4$), the cost tends to be overestimated; however, as the lead-time variability increases from low ($L_{C_v} = 0.4$) to medium ($L_{C_v} = 1.0$) the cost will be less overestimated, and eventually become underestimated when the coefficient of lead-time is high ($L_{C_v} = 1.6$). On the other hand, while the demand coefficient of variation is high ($D_{C_v} = 1.6$) the cost is always underestimated and the relative error increases as the lead-time coefficient increases from low to high.




Figure 3. Interaction Plot of D_{C_v} vs. L_{C_v} of Cost Relative Error.

Based on the observations of the summary statistics and ANOVA results, it is apparent that the model tested in this study does not perform well when the assumptions are violated (Poisson demand to non-Poisson demand with C_v of 0.4 and 1.6; constant lead-time to non-constant lead-time with C_v of 0.4, 1.0, and 1.6). When the demand variability is low ($D_{C_v} = 0.4$), less backorders will occur in the actual system, i.e. the model overestimates the number of backorders and hence the total cost. The model continues to overestimate the backorders and total cost until the lead-time variability is high enough (at $L_{C_v} = 1.6$) to increase the number of backorders in the system. The high lead-time variability increases the uncertainty, and hence more chances of backordering within the system. Nevertheless, while the demand variability is high ($D_{C_v} = 1.6$), the system encounters higher uncertainty that more backorders arise, i.e. the model underestimates the backorders and total cost. As the lead-time variability increases, more and more backorders happen within the system, and the model prediction gets worse. It is not surprising to find out that the maximum cost relative error (the most overestimated) in the experiments comes from the conditions of low demand and lead time variability ($D_{C_v} = 0.4$, $L_{C_v} = 0.4$); the high demand and lead-time variability ($D_{C_v} = 1.6$, $L_{C_v} = 1.6$) correspond to the minimum cost relative error (the most underestimated).

The experimental results also indicate that the demand factor (D) has considerable impact on the model's predictions. When the demand is low ($D = 0.1$), the model tends to give less error when estimating the cost. When the demand is high ($D = 1.0$), more fluctuations are observed. Higher demand provides more opportunities for uncertainty (causing error) and hence the model estimation is poorer. Besides the demand factor, other factors such as the retailer batch order quantity

(Q_r), retailer number (N) and backorder cost factor (P_r) also influence the model prediction in some ways, but no particular important trend was identified.

5. CONCLUSIONS AND FUTURE RESEARCH

This study evaluates the behavior of a (R, Q) multi-echelon inventory model in estimating the total system cost under demand and lead-time conditions differing from those originally assumed. The tested conditions are stationary non-Poisson demand and random lead-time.

Neither overestimation nor underestimation is good for a company who uses these analytical models. If the total system cost is overestimated, the company might hold too much capital for its distribution system; capital that could be invested elsewhere. If the total system cost is underestimated, the company might not have enough capital to cope with the actual situation. Underestimation is considered to cause more serious losses to the company because the company will need to spend more money, and hence the profits are reduced. The simulation results of this study show that there was significant error in estimating the total system cost (tends to be underestimated) when the model assumptions are violated. The cost relative error ranged between 34% and -40%.

A significant factor was identified from this study: the demand rate. When the demand is low, the model will give less error in estimating the total cost. Thus, the model is more suitable to be used for low demand items. Under the conditions that have more demand variation and uncertainty than the model assumed, the model will recommend a policy of carrying less safety stock than is needed, and hence more backorders occur and the total system cost is increased. If the actual demand is less variable, less backorders will happen under the recommended inventory policy, which means the model overestimates and influences the system to carry more inventory.

When considering the lead-time variability, the amount of error increases as the variance of the lead-time distribution increases. Higher variability in the lead-time will require higher safety stock to maintain the same service level. If the model neglects the effect of lead-time variability, more backorders will arise from the recommended inventory policy. Therefore, the company can gain benefit by reducing the lead-time variability, such as choosing more consistent transportation providers.

As a conclusion, the (R, Q) two-echelon inventory model considered in this study is not robust enough to provide good estimation for the total system cost under the considered testing conditions. Careful consideration is needed before the usage of these models and the implementation of any recommended inventory policy. We recommend that practitioners evaluate the potential use of multi-echelon methodologies in a simulation study before implementation.

Several extensions to this research are possible. The inventory policy recommended by an analytical model depends upon the tradeoff between the cost components under certain assumptions. Therefore, the estimation of the components of the total cost should be examined carefully to investigate the relationship of cost and its components. One can also perform the analysis on a broader range of model parameters, such as lead-times of greater than one day, and larger holding cost factors. Lead-times of higher than one day are more realistic in some situations; however, more variability in the system will be observed. Besides, one can also consider the non-identical retailers case. In the real world, it is very unlikely that the retailers are completely identical in demand, order quantity, or lead-time. The Axsater (2000) model can evaluate this case; however, the most critical issue is how to define the non-identical test cases and still be able to perform a fair comparison on the models. The results of this study show that factors other than the demand factor significantly influence the model, but another study is needed to identify the important trends and model behaviors.

The results of this research also provide the following future research directions:

- Test the model by using non-stationary Poisson demand process. Most companies experience non-stationary demand, and some might have a forecasting system that can capture such demand. We have analyzed the model with a piecewise non-stationary Poisson demand process. A paper is in preparation for the publication of this work.
- Investigate to develop a more robust model, especially in estimating the average backorder level because it is the major factor in determining the cost. This requires more detailed study of the lead-time demand distributions to find a more robust form for the representation of the system. Heuristic approaches can also be used to reduce the complexity of computational efforts.

6. REFERENCES

1. Axsater, Sven. (1990). Simple Solution Procedures for a Class of Two-Echelon Inventory Problems. Operation Research, 38(1): 64-69.
2. Axsater, Sven. (1993a). Continuous Review Policies for Multi-Level Inventory Systems with Stochastic Demand. Chapter 4 in S.C. Graves et al. (eds.), Handbooks in OR and MS, 4: 175-197.
3. Axsater, Sven. (1993b). Exact and Approximate Evaluation of Batch-Ordering Policies for Two-Level Inventory Systems. Operation Research, 41(4): 777-785.
4. Axsater, Sven. (1995). Approximate Evaluation of Batch-Ordering Policies for a One-Warehouse, N Non-Identical Retailer System under Compound Poisson Demand. Naval Research Logistics, 42: 807-819.
5. Axsater, Sven. (1998). Evaluation of Installation Stock Based (R, Q)-Policies for Two-Level Inventory Systems with Poisson Demand. Operation Research, 46(3): S135-S145.
6. Axsater, Sven. (2000). Exact Analysis of Continuous Review (R, Q) Policies in Two-Echelon Inventory Systems with Compound Poisson Demand. Operation Research, 48(5): 686-696.
7. Bagchi, Uttaraya, Jack C. Hayya, and J. Keith Ord. (1984). Concepts, Theory, and Techniques: Modeling Demand During Lead Time. Decision Sciences, 15: 157-176.
8. Banks, Jerry, and J. P. Spoerer. (1986). Inventory Policy for the Continuous Review Case: A Simulation Approach. Annals of the Society of Logistics Engineering, 1(1): 51-65.
9. Cohen, M. A., Kamesam, Kleindorfer, Lee, and Tekerian. (1990). Optimizer: IBM's Multi-Echelon Inventory System for Managing Service Logistics. Interfaces, 20(1): 65-82.
10. Cohen, M. A., P.R. Kleindorfer, and H.L. Lee. (1986). Optimal Stocking Policies for Low Usage Items in Multi-Echelon Inventory System. Naval Research Logistics Quarterly, 33: 17-38.
11. Deurmeyer, B., and L. B. Schwarz. (1981). A Model for the Analysis of System Service Level in Warehouse/ Retailer Distribution Systems: the Identical Retailer Case. In L. B. Schwarz (eds.), Studies in Management Sciences, Vol. 16, Multi-Level Production/ Inventory Control Systems, North-Holland, Amsterdam, 163-193.
12. Forsberg, Rolf. (1996). Exact Evaluation of (R, Q)-policies for Two-level Inventory Systems with Poisson demand. European Journal of Operational Research, 96: 130-138.
13. Forsberg, Rolf. (1997). Evaluation of (R, Q)-policies for Two-level Inventory Systems with Generally Distributed Customer Inter-arrival Times. European Journal of Operational Research, 99: 4001-411.
14. Fortuin, Leonard. (1980). Five Popular Probability Density Functions: A Comparison in the Field of Stock-Control Models. Journal of the Operational Research Society, 31: 937-942.
15. Fotopoulos, Stergios, Min-Chiang Wang, and S. Subba Rao. (1988). Safety Stock Determination with Correlated Demands and Arbitrary Lead Times. European Journal of Operational Research, 35: 172-181.
16. Hopp, Wallace J. and Mark L. Spearman. (2000). Factory Physics: Foundations of Manufacturing Management, 2nd edition, Irwin/McGraw-Hill, New York.
17. Kelton, David, Randall Sadowski, and Deborah Sadowski. (1998). Simulation with Arena, McGraw-Hill Inc., Boston.
18. Lagodimos, A. G., A. G. De Kok, and J. H. C. M. Verrijdt. (1995). The Robustness of Multi-echelon Service Models under Autocorrelated Demands. Journal of the Operational Research Society, 46(1): 92-103.
19. Lau, Hon-Shiang, and Ahmed Zaki. (1982). The Sensitivity of Inventory Decisions to the Shape of Lead Time Demand Distribution. IIE Transactions, 14(4): 265-271.
20. Law, Averill M., and W. David Kelton. (2000). Simulation Modeling and Analysis, 3rd edition, McGraw-Hill Inc., New York.
21. Moinszadeh, K., and H. L. Lee. (1986). Batch Size and Stocking Levels in Multi-echelon Repairable Systems. Management Sciences, 32(12): 1567-82.
22. Muckstadt, J. A. (1973). A Model for a Multi-Item, Multi-Echelon, Multi-Indenture Inventory System. Management Sciences, 20: 472-481.
23. Naddor, Eliezer. (1978). Sensitivity to Distributions in Inventory Systems. Management Sciences, 24(16): 1769-1772.
24. Needham, Paul M., and Philip T. Evers. (1998). The Influence of Individual Cost Factors on the Use of Emergency Transshipments. Transportation Research Part E: Logistics and Review, 34(2): 149-160.
25. Sherbrooke, C. C. (1968). METRIC: A Multi-Echelon Technique for Recoverable Item Control. Operations Research, 16: 122-141.
26. Silver, Edward A, David F. Pyke, and Rein Peterson. (1998). Inventory Management and Production Planning and Scheduling, 3rd edition, John Wiley & Sons, New York.
27. Svoronos, Antony, and Paul Zipkin. (1988). Estimating the Performance of Multi-level Inventory Systems. Operation Research, 36(1): 57-72.

Proceedings of the 6th Annual
International Conference on Industrial
Engineering – Theory, Applications, and
Practice, San Francisco, CA, USA,
November 18-20, 2001

28. Tyworth, John E., and Liam O’Neill. (1997). Robustness of the Normal Approximation of Lead-time Demand in a Distribution Setting. Naval Research Logistics Quarterly, 44(2): 165-86.
29. Zipkin, Paul H. (1986). Stochastic Leadtimes in Continuous-Time Inventory Models. Naval Research Logistics Quarterly, 33: 763-774.
30. Zipkin, Paul H. (2000). Foundations of Inventory Management, McGraw-Hill, New York.