

**EVALUATING VARIANCE REDUCTION TECHNIQUES WITHIN
A SAMPLE AVERAGE APPROXIMATION METHOD FOR A
CONSTRAINED INVENTORY POLICY OPTIMIZATION PROBLEM**

Yasin Ünlü
Manuel D. Rossetti

The University of Arkansas
4207 Bell Engineering Center
Fayetteville, AR 72701, USA

ABSTRACT

This paper examines a constrained stochastic inventory optimization problem by means of sample average approximations (SAA). The problem is formulated based on the lead time demand parameters. Lead time demands are sampled by a bootstrap method that is performed by randomly generating demand values over deterministic lead time values. In order to increase the efficiency of solving an SAA replication, a number of variance reduction techniques (VRT) are proposed, namely: antithetic variates, common random numbers and Latin hypercube sampling methods. A set of experiments investigates the quality of these VRTs on the estimated optimality gap and gap variance results for different demand processes. The results indicate that the use of VRTs produces significant improvements over the crude Monte Carlo sampling method on all test cases.

1 INTRODUCTION

In this paper, a constrained stochastic inventory optimization problem is examined within a sample average approximation (SAA) method. The problem is defined as the minimization of the total expected relevant inventory costs subject to a service level constraint. The foregoing problem arises in the context of the continuous review (r, q) inventory system where the optimal policy parameters are determined by discrete values of r and q . We consider the ready rate as the underlying service level constraint. The ready rate is known as the fraction of time with positive stock on-hand (Axsäter 2006, p94). Even though the problem contains a service level constraint, we propose an approach to represent it in a single objective function, which facilitates the application of the SAA method.

As far as the joint optimization of the discrete policy parameters is concerned, the defined constrained problem is known as a hard problem (Zipkin 2000, p226). A heuristic algorithm is proposed by (Zipkin 2000, p227) that states that the policy obtained through the heuristic does not guarantee the minimization of the true costs. His heuristic is developed based on the algorithm proposed by (Federgruen and Zheng 1992) that is focused on the unconstrained problem. It should be noted that the algorithm proposed by (Federgruen and Zheng 1992) can also be used to obtain exact solutions for the constrained problem only for cases where the lead time demand (LTD) is Poisson. For general LTD cases, (Agrawal and Seshadri 2000) propose an optimization algorithm which is developed based on the search in the regions defined by the distribution free bounds on the policy parameter q . However, their algorithm is still dependent upon a given LTD distribution due to the calculation of the total relevant inventory costs.

In most practices, the LTD distribution is not known or unavailable in a closed mathematical form. Therefore, the analytical formulas to exactly calculate the total inventory costs can not be directly used.

The solution of the inventory optimization problem whose cost function is represented in the expected form is optimal in the expected sense. The calculation of expected costs is expensive since it typically involves enumerating all possible LTD outcomes which may be observed as combinatorially many scenarios depending on the input parameters. We apply the SAA method due to (Ahmed and Shapiro 2002) and (Kleywegt and Shapiro 2001) in order to estimate the expected costs in the problem. The motivation behind employing SAA is the theoretical fact that the solution to the approximation problem exponentially converges to the optimal solution as the number of scenarios increases (Kleywegt and Shapiro 2001). The SAA method has been applied to many different stochastic problem domains. Examples include the stochastic bidding and stochastic scheduling problems (Greenwald, Guillemette, Naroditskiy, and Tschantz 2006), vehicle assignment, aircraft allocation, network design and cargo flight scheduling (Linderoth, Shapiro, and Wright 2006).

The optimization procedure described in this study requires no explicit LTD model. However, the problem formulation is still dependent upon the LTD parameters. The optimization procedure is predicated on evaluating each candidate solution which is associated with each given q whose set is determined by applying the distribution free bounds proposed by (Agrawal and Seshadri 2000). The candidate solution is generated by solving the SAA problem which is constructed based on the sample of the LTD values. A bootstrapping method is applied to generate LTD values from a given demand distribution over a given lead time value. We apply SAA statistical techniques to evaluate each candidate solution. This evaluation relies on the specified precision on the optimization gap and gap variance, which are subject to vary depending on the sample size. Typically a large sample size results in more precision while it often causes an increase in the computational time. One can apply a variance reduction technique which helps reduce the required sample size to reach a specified precision in the gap variance value. It should be noted that the variance observed during the evaluation procedure also effects the precision of the optimization gap. Thus, identifying a variance reduction technique can be considered a key to a potential efficient SAA based optimization algorithm.

The focus of this paper is to evaluate a number of variance reduction techniques which may be applicable in the design of an efficient optimization algorithm. In this respect, three variance reduction techniques; namely, antithetic variates (AV), Latin hypercube sampling method (LHS) and common random numbers (CRN) are considered in this study. The AV, LHS and CRN are the techniques that are often recommended for reducing variance (Law and Kelton 1999). The variance reduction observed under these VRTs is compared with results observed through the crude Monte Carlo sampling method. This paper discusses how these VRTs can be utilized in the context of the constrained stochastic inventory optimization problem. The details of the problem will be introduced in the next section. Section 3 describes the SAA based solution procedure and the variance reduction techniques. Section 4 discusses the effectiveness of the underlying variance reduction techniques. Finally, Section 5 provides concluding remarks and directions for the future research.

2 PROBLEM FORMULATION

The constrained optimization problem arises in the following inventory system. The inventory is reviewed continuously at a single stage for a single item and controlled by the (r, q) policy with the following mechanism. Whenever the inventory position IP drops to or below an integer value of re-order level (r) , the amount of q units of items is issued to replenish the inventory, and the ordered items arrive at the inventory system after a given constant time delay L . Let $E[D]$ be the expected value of the demand. The lead times are assumed to be independent of the demand process. Let Y be the total demand during lead time, which is a random variable with the expected value $E[D]L$.

A bootstrapping method is of interest in generating lead time demand values which will be utilized in the SAA procedure. Since the lead times are assumed to be independent of the demand process, the following procedure can be used to generate Y . In most practices, even if the inventory system faces a

compound demand process, the information related to the demand during a unit time (per day, per week, etc.) is available. By using the given mean (μ_X) and variance (σ_X^2) parameter information of the demand during a unit time, demand amounts are generated for discrete unit time points (t). The total sum value of the demand observed over a fixed lead time (L) gives independent and identically distributed Y values. This procedure is independently repeated to produce a sample of Y . We generate the unit time demand component from a distribution. Let X_t be the generated demand amount during a time unit and $X_t \sim f_X(\cdot)$, then clearly,

$$Y = \sum_{t=1}^L X_t \tag{1}$$

The inventory system faces a discrete compound demand process. That is, customer demands arrive at any point in time with an amount of items as a batch or as individual units. We assume that the inventory system processes the demand in individual units even if it arrives as a batch of items. The reader is referred to (Rossetti and Ünlü 2011) for details of the inventory system. Based on the described inventory system, the following exact formulations are used to compute the inventory performance measures of the ready rate ($E[RR]$), the expected number of inventory on-hand levels ($E[I]$) and the expected order frequency ($E[OF]$) (Zipkin 2000).

$$E[RR] = 1 - \frac{1}{q} [G^1(r) - G^1(r+q)] \tag{2}$$

$$E[I] = \frac{1}{2}(q+1) + r - Y + \frac{1}{q} [G^2(r) - G^2(r+q)] \tag{3}$$

$$E[OF] = \frac{E[D]}{q} \tag{4}$$

where $E[D] = E[Y]/L$, $G_F^1(\cdot)$ and $G_F^2(\cdot)$ are the first and second order loss functions of the lead time demand distribution F . Let $[y]^+$ be denoted $\max\{0, y\}$. Then it follows that $G_F^1(y) = E[[Y - y]^+]$ and $G_F^2(r) = 0.5E[[Y - r]^+ [Y - r - 1]^+]$.

In modeling the inventory optimization problem, a cost structure is imposed on $E[OF]$ and $E[I]$. The backordered demand is controlled by imposing a service level constraint. The ready rate is used as the service level in this study due to the existence of a tractable analytical formulation. The policy optimization of the continuous review (r, q) system is performed by solving the corresponding stochastic inventory problem. The goal is to obtain the optimal discrete policy parameters r and q which minimize the sum of ordering and holding costs subject to the constraint that the ready rate should be at least equal to γ . Let k be the fixed cost to place an order and h be the holding cost per unit per unit time. The cost measures are assumed to be positive in order for q to be finite positive integer value while policy parameter r takes any finite integer values greater than or equal to $-q$. For a pair of (r, q) , denote the expected total cost $E[TC(r, q)]$. Then the optimization problem is given as follows.

Optimization Problem P1:

$$\min E[TC(r, q)] = kE[OF] + hE[I] \tag{5}$$

$$\text{subject to } E[RR] \geq \gamma \tag{6}$$

The optimization problem can be represented as

$$\begin{aligned} \min E [TC(r, q)] = & E \left[\frac{kE [D]}{q} + h \left(\frac{q+1}{2} + r - Y \right) \right. \\ & + \frac{h}{2q} \left(\max \{Y - r, 0\} \max \{Y - r - 1, 0\} \right. \\ & \left. \left. - \max \{Y - (r + q), 0\} \max \{Y - (r + q) - 1, 0\} \right) \right] \end{aligned} \quad (7)$$

$$\text{subject to } 1 - \frac{1}{q} E [\max \{Y - r, 0\} - \max \{Y - (r + q), 0\}] \geq \gamma \quad (8)$$

For a given q , let $r(q)$ be the re-order point that satisfies (8). Let $S = \{r : r \geq r(q)\}$. Thus, S represents the set of possible r that satisfy the service level. For a given q , let $r^*(q)$ be the optimal re-order point to the optimization problem defined above. That is, $r^*(q) = \arg \min_{r \in S} E [TC(r, q)]$. Then the optimization problem can be rewritten as a single objective function as follows:

Optimization Problem P2:

$$\begin{aligned} \min E [TC(q, r^*(q))] = & E \left[\frac{kE [D]}{q} + h \left(\frac{q+1}{2} + r^*(q) - Y \right) \right. \\ & + \frac{h}{2q} \left(\max \{Y - r^*(q), 0\} \max \{Y - r^*(q) - 1, 0\} \right. \\ & \left. \left. - \max \{Y - (r^*(q) + q), 0\} \max \{Y - (r^*(q) + q) - 1, 0\} \right) \right] \end{aligned} \quad (9)$$

If Y follows the pure Poisson distribution (i.e., demand size is 1 and inter-arrivals are exponentially distributed), then (9) is convex since $r(q)$ is unimodal. In this case, P1 can be solved through the algorithm proposed by (Federgruen and Zheng 1992) with the help of Lagrange multipliers. Unfortunately (9) is generally not convex. The search of the optimal policy requires a full enumeration of q if Y follows a distribution other than Poisson (Zipkin 2000, p226). In the case where Y follows a known distribution (e.g., gamma), (Agrawal and Seshadri 2000) propose an optimization algorithm for P1. This paper, on the other hand, is focused on the cases where the lead time demand distribution model is not known (or not available in the closed mathematical form). The random lead time demand variable Y is bootstrapped by generating random demand values over a lead time. We perform the full enumeration over a finite set of q which is determined by bounds applied on q (Agrawal and Seshadri 2000). Each possible q value creates a candidate solution. These candidate solutions are evaluated and the best solution is selected from the set. We evaluate candidate solutions by using the sample average approximation (SAA) technique, which allows the estimation of the expected value in (9). The next section gives the details related to the solution procedure including the SAA method and the generation of Y .

3 SOLUTION PROCEDURE

The SAA technique is applied in this paper as follows. First, P candidate solutions are determined. Each candidate solution refers to a q and $r(q)$ from a finite set. As mentioned earlier, this set can be obtained by applying the distribution free bounds on q proposed by (Agrawal and Seshadri 2000). Each candidate

solution is evaluated by SAA statistical tools to see their quality viewed as a candidate for solving the true problem. Then the best solution is selected from the set of the evaluated candidate solutions. The next section discusses the evaluation procedure in the context of optimization gap and variance reduction.

3.1 Optimization Gap Construction and Variance Reduction

The evaluation procedure involves constructing the optimization gap for each candidate solution \hat{x} . The candidate solutions (i.e., pair of $(q, r^*(q))$) are obtained based on the following fact. For any given $q \geq 1$, the value of $\hat{q}_N(x) = \frac{1}{N} \sum_{i=1}^N TC(x, Y^i)$ increases in r where $Y^i \in F_{LTD}(\cdot)$. In addition, there exists $r^*(q) \in S$ such that $r^*(q) = \arg \min_{r \in S} \hat{q}_N(x)$. Thus, for a fixed q , the optimal value of (9) can be obtained by the minimum feasible value of r . This will provide $r^*(q)$. The solution is performed satisfying the constraint (8) in the sample average sense. That is, for a given q and a set of Y , $r^*(q)$ is the minimum value of r that satisfies the following:

$$1 - \frac{1}{q} \left\{ \frac{1}{N} \left(\sum_{i=1}^N [\max \{Y^i - r, 0\} - \max \{Y^i - (r+q), 0\}] \right) \right\} \geq \gamma \tag{10}$$

For a given value of q , it is trivial to obtain $r^*(q)$ through a simple line search. The candidate solution is denoted by $\hat{x} = (q, r^*(q))$. Clearly, the pair $(q, r^*(q))$ minimizes the approximation $\hat{q}_N(\hat{x})$. This refers to the fact that the candidate solution $\hat{x} = (q, r^*(q))$ is generated by “solving the corresponding SAA problem to optimality.” However, this solution should be evaluated to see its quality viewed as a candidate for solving the true problem. Notice that the true objective function value of this solution is different from the approximated one. For a given q , we apply statistical methods to estimate bounds for the true objective value. In what follows, we discuss the details of constructing the optimization gap based on upper and lower bounds.

3.1.1 Independent Sampling (CMC: Crude Monte Carlo Method)

Upper Bound Estimation: In order to estimate an upper bound, we first construct an unbiased estimator of $g(\hat{x})$. We generate M independent and identically distributed (i.i.d.) batches of random samples. In addition, each batch consists of N random elements (i.e., the sample of Y). Let the generated batches be denoted by i.i.d. random elements of $w^{1,j}, w^{2,j}, \dots, w^{N,j}$ $j = 1, 2, \dots, M$. Then, the unbiased property of each batch can be denoted by

$$E \left[\hat{q}_N^j(\hat{x}) = \frac{1}{N} \sum_{i=1}^N TC(\hat{x}, w^{i,j}) \right] = g(\hat{x}) \tag{11}$$

Let $\overline{g(\hat{x})}_{N,M}$ be an estimate for $g(\hat{x})$. Then

$$\overline{g(\hat{x})}_{N,M} = \frac{1}{M} \sum_{j=1}^M [\hat{q}^j(\hat{x})] \tag{12}$$

is an unbiased estimate of $g(\hat{x})$. The associated sample variance estimator is obtained by

$$\hat{\sigma}_{UB}^2 = \frac{1}{M-1} \sum_{j=1}^M \left[\hat{q}^j(\hat{x}) - \overline{g(\hat{x})}_{N,M} \right]^2 \tag{13}$$

Since $g(\hat{x}) \geq v^*$ (for minimization type problems), an approximate $100(1-\alpha)\%$ upper bound estimate is given by

$$\hat{UB}_{N,M} = \overline{g(\hat{x})}_{N,M} + \frac{t_{\alpha,v} \sigma_U}{\sqrt{M}} \tag{14}$$

where $v = M-1$ and $t_{\alpha,v}$ is the α -critical value of the t-distribution with v degrees of freedom.

Lower Bound Estimation: Let $\hat{v}_{N'}$ be the optimal value of an SAA problem based on sample size N' . Then, a lower bound can be estimated based on the fact that $v^* \geq E[\hat{v}_{N'}]$. We estimate a lower bound for $E[\hat{v}_{N'}]$. $E[\hat{v}_{N'}]$ can be estimated by solving SAA problems several times and averaging the calculated optimal values. M' SAA problems are created based on generated i.i.d. batches which contain i.i.d. random elements of $w^{1,j}, w^{2,j}, \dots, w^{N',j}$ $j = 1, 2, \dots, M'$. Let $\hat{v}_{N'}^1, \hat{v}_{N'}^2, \dots, \hat{v}_{N'}^{M'}$ be the computed optimal values of the SAA problems. Then

$$\bar{v}_{N',M'} = \frac{1}{M'} \sum_{i=1}^{M'} \hat{v}_{N'}^i \tag{15}$$

is an unbiased estimator of $E[\hat{v}_{N'}]$. We can estimate the variance of $\bar{v}_{N',M'}$ as follows

$$\hat{\sigma}_{LB}^2 = \frac{1}{(M' - 1)} \sum_{i=1}^{M'} (\bar{v}_{N',M'} - \hat{v}_{N'}^i)^2 \tag{16}$$

An approximate $100(1-\alpha)\%$ lower bound for $E[\hat{v}_{N'}]$ is then given by

$$\hat{LB}_{N',M'} = \bar{v}_{N',M'} - t_{\alpha,v} \hat{\sigma}_{LB} \tag{17}$$

The quality of \hat{x} can be measured by the optimality gap

$$gap(\hat{x}) = g(\hat{x}) - v^* \tag{18}$$

We outline a statistical procedure for estimating this optimality gap via upper bound (\hat{UB}) and lower bound (\hat{LB}) analysis. Thus,

$$gap(\hat{x}) = \hat{UB} - \hat{LB} \tag{19}$$

and

$$\sigma_{gap}^2 = \hat{\sigma}_{UB}^2 + \hat{\sigma}_{LB}^2 \tag{20}$$

We determine the candidate solution (\hat{x}) during the upper bound estimation procedure. The candidate solution, the optimality gap and gap variance are subject to change according to different sample sizes. Although a large sample size gives a better estimate, it increases the computational time of the evaluation procedure. Therefore, the SAA parameters N (batch size for UB estimation), N' (batch size for LB estimation), M (number of batches for UB estimation) and M' (number of batches for LB estimation) should be wisely determined in the optimization algorithm development phase. The optimality gap and gap variance are considered as major precision criteria in the development of an optimization procedure that evaluates the candidate solutions. It should be noted that the variance observed during the evaluation procedure also effects the precision of the optimization gap. In this respect, reducing variance is the key to an efficient SAA based optimization algorithm.

Note that the random elements within a batch need not be i.i.d provided that the statistical bounds are constructed based on the i.i.d. batches. The reader is referred to (Mak, Morton, and Wood 1999) where the underlying theory is discussed in detail. We now present a number of variance reduction techniques by utilizing the foregoing theory.

3.1.2 Antithetic Variates (AV)

In independent sampling method, $\overline{g(\hat{x})}_{N,M}$ (or $\bar{v}_{N',M'}$) is obtained based on M (or M') independent batches, respectively. As far as antithetic variates are concerned, independent batches can be generated as follows. We first find an estimate for $g(\hat{x})$ (or $E[\hat{v}_{N'}]$) based on a batch of random sample of size N (or N'). Next, we find another estimate for $g(\hat{x})$ (or $E[\hat{v}_{N'}]$) based on the same sample size. However, the second batch

contains the antithetics of the first batch. In order to reach an independent estimate of $g(\hat{x})$ (or $E[\hat{v}_N]$), the average of those two estimates is obtained. Therefore, $\overline{g(\hat{x})}_{N,M}$ (or $\bar{v}_{N',M'}$) is obtained based on $M/2$ (or $M'/2$) i.i.d. batches of random samples. Then the same expressions are used to obtain an optimality gap for the candidate solution. The AV is one of the most applied variance reduction techniques. An application of the AV to the newsvendor problem is studied by (Freimer, Thomas, and Linderoth 2010).

3.1.3 Latin Hypercube Sampling (LHS)

We generate random elements in each batch of samples via the Latin hypercube sampling method. If the lead time is given as 1 time unit (i.e., $L = 1$), then in this one-dimensional sampling, we divide the interval $[0, 1]$ into N (sample size) equal segments. The lead time demand value is drawn uniformly from the i^{th} segment. That is, the lead time demand value under LHS is uniformly distributed on $[(i - 1)/N, i/N]$. If the lead time is greater than 1 unit, then in this multi-dimensional sampling, the range of $[0, 1]$ is portioned into N non-overlapping intervals of equal probability $1/N$. From each interval one demand value is selected randomly according to the probability density of the interval. The N values of D_1 are paired in a random manner with values of D_2 , these pairs are then paired similarly with values of D_3 and so on, until N samples of L time units are formed. Then the corresponding lead time demand value (Y) is obtained by the total sum value of the demand observed over L time units (i.e., $Y = D_1 + D_2 + \dots + D_L$). The reader is referred to (Matala 2008) for the accuracy of LHS method and the simple strategy to evaluate N for general problem domains.

3.1.4 Common Random Numbers (CRN)

The common random numbers method within the sample average approximation is proposed by (Mak, Morton, and Wood 1999). The idea is to use the batch means approach to directly estimate the optimization gap for the candidate solution. Based on a batch of sample of size N , an estimate for $gap(\hat{x})$ can be obtained by

$$E \left[\frac{1}{N} \sum_{i=1}^N TC(\hat{x}, w^i) - \min_{x \in \mathcal{X}} \frac{1}{N} \sum_{i=1}^N TC(x, w^i) \right] \tag{21}$$

Note that the upper and lower bounds are estimated by using the same batch, which can be considered as an application of the common random numbers. The optimization gap and gap variance can be estimated based on separately estimated upper and lower bounds by using the formulas introduced in Section 3.1.1.

4 EXPERIMENTAL STUDY

The quality of the variance reduction techniques is computationally investigated on the estimated optimality gap and gap variance results across a large set of test cases. The results are collected under different demand models; namely, Poisson, negative binomial and gamma. The test cases are generated based on the combination of the low and high values of a number of experimental . These factors are given in Table 1.

Table 1: Experimental Factors.

Level	Target Service Level	Lead Time	Mean LTD	Variance LTD	Ordering Cost	Holding Cost
Low	0.90	1	1.8	4	50	1
High	0.95	4	3.6	8	100	10

Based on the given experimental factors, the algorithm proposed by (Agrawal and Seshadri 2000) indicates that the optimal reorder quantity can take values over the range between 1 and 30. Therefore,

test cases are generated based on the combination of the given values from Table 1 and reorder quantity enumerated over the set $\{1, 2, 3, \dots, 30\}$. This creates $(2)2(2)2(2)2(30) = 1920$ different test cases. For each test case, the candidate solution (i.e., $\hat{x} = (q, r^*(q))$) is obtained based on a given sample size. Then an estimate is obtained for the upper bound, lower bound, (accordingly, optimization gap on the candidate solution) and their variances estimates. The average values of these estimates are tabulated in Table 2 and Table 3 for different total sample size values. Notice that tables show blank cells for upper and lower bounds of the sampling technique CRN whose optimization gap and gap variance are directly estimated.

Table 2: Optimization Gap and Variance Results for Total Sampling 1000.

	Sampling Technique	Upper Bound			Lower Bound			Total Sampling	Optimization Gap	Total		Variance Reduction
		<i>N</i>	<i>M</i>	UB Variance	<i>N'</i>	<i>M'</i>	LB Variance			Gap	Variance	
Pure Poisson	CMC	50	10	0.0302	50	10	0.4721	1000	2.8488	0.5023	-	
	AV	50	5	0.0399	50	5	0.0881	1000	2.5540	0.1280	75%	
	CRN	100	10	-	-	-	-	1000	1.9162	0.0219	96%	
	LHS	50	10	0.0006	50	10	0.2178	1000	1.0169	0.2184	57%	
Negative Binomial	CMC	50	10	0.0559	50	10	1.6389	1000	5.8829	1.6948	-	
	AV	50	5	0.0618	50	5	0.2615	1000	4.0354	0.3233	81%	
	CRN	100	10	-	-	-	-	1000	3.1389	0.0501	97%	
	LHS	50	10	0.0013	50	10	0.5932	1000	2.1755	0.5945	65%	
Gamma	CMC	50	10	0.0529	50	10	1.9523	1000	6.4134	2.0052	-	
	AV	50	5	0.0508	50	5	0.2930	1000	4.3434	0.3437	83%	
	CRN	100	10	-	-	-	-	1000	3.2256	0.0448	98%	
	LHS	50	10	0.0011	50	10	0.6269	1000	2.2665	0.6280	69%	

Table 3: Optimization Gap and Variance Results for Total Sampling 2000.

	Sampling Technique	Upper Bound			Lower Bound			Total Sampling	Optimization Gap	Total		Variance Reduction
		<i>N</i>	<i>M</i>	UB Variance	<i>N'</i>	<i>M'</i>	LB Variance			Gap	Variance	
Pure Poisson	CMC	100	10	0.0156	100	10	0.3843	2000	2.2486	0.3999	-	
	AV	100	5	0.0127	100	5	0.0668	2000	1.8927	0.0795	80%	
	CRN	200	10	-	-	-	-	2000	1.4477	0.0173	96%	
	LHS	100	10	0.0002	100	10	0.1392	2000	0.6324	0.1393	65%	
Negative Binomial	CMC	100	10	0.0282	100	10	0.9693	2000	4.1098	0.9975	-	
	AV	100	5	0.0184	100	5	0.1655	2000	3.1790	0.1839	82%	
	CRN	200	10	-	-	-	-	2000	2.2823	0.0265	97%	
	LHS	100	10	0.0004	100	10	0.2869	2000	1.2596	0.2873	71%	
Gamma	CMC	100	10	0.0267	100	10	1.1145	2000	4.4636	1.1413	-	
	AV	100	5	0.0178	100	5	0.1893	2000	3.3281	0.2071	82%	
	CRN	200	10	-	-	-	-	2000	2.3415	0.0258	98%	
	LHS	100	10	0.0003	100	10	0.2769	2000	1.1935	0.2772	76%	

As can be noted from both tables, the VRTs are effective in terms of reducing the optimization gap variance estimated through the crude Monte Carlo sampling method. In addition, the value of the optimization gap is inclined to be smaller under the applied VRTs. For example, as can be seen from Table 1, for the negative binomial model with sample size 1000, CMC sampling method yields 5.8829 for optimization gap and 1.6948 for optimization gap variance while the optimization gap and gap variance under CRN are observed 3.1389 and 0.0501, respectively. CRN is able to reduce the optimization gap variance by 97%. The minimum optimization gap value is always estimated by the Latin hypercube sampling

method. The common random numbers yield the minimum estimated optimization gap results among all considered sampling methods. By comparing the observed variance results in Table 2 and Table 3, one can note that the VRTs are more effective under a larger amount of sampling since VRTs are able to reduce the observed variance by similar percentages. For example, for the gamma distribution CRN is able to reduce the observed optimization gap variance by approximately 98% for each sample size 1000 and 2000. Therefore, VRTs yield much smaller variance values for large sample sizes.

The statistical package MINITAB provides the multiple comparison procedure referred to as “Tukey-Kramer HSD” ((Tukey 1953) and (Kramer 1956)) which facilitates analyzing the overall performance of the VRTs. Based on the least square means for each pair of the VRTs, the results are tabulated in a categorized manner. Each category is represented by a letter in a column. Table 4, Table 5 and Table 6 tabulate the results of the procedure for the Poisson, negative binomial and gamma models, respectively with 95% confidence level and with sample size 2000. If two sampling techniques share a letter, then they are regarded as not significantly different from each other. In this respect, the performance if the VRTs is sorted in descending order as follows: $\{CRN\} > \{AV\} > \{LHS\} > \{CMC\}$ for the Poisson and negative binomial models and $\{CRN\} > \{AV, LHS\} > \{CMC\}$ for the gamma model. For all models, *AV*, *LHS* are the two VRTs whose performances are significantly higher than *CMC* while the performance of *CRN* is significantly higher than other VRTs.

Table 4: Comparisons for all Sampling Techniques using Tukey-Kramer HSD for the Poisson Model.

Models	Category 1	Category 2	Category 3	Category 4	Mean
CMC	D				0.3999
AV		B			0.0795
CRN			A		0.0173
LHS				C	0.1393

Table 5: Comparisons for all Sampling Techniques using Tukey-Kramer HSD for the Negative Binomial Model.

Models	Category 1	Category 2	Category 3	Category 4	Mean
CMC	D				0.9975
AV		B			0.1839
CRN			A		0.0265
LHS				C	0.2873

Table 6: Comparisons for all Sampling Techniques using Tukey-Kramer HSD for the Gamma Model.

Models	Category 1	Category 2	Category 3	Category 4	Mean
CMC	D				1.1413
AV		B			0.2071
CRN			A		0.0258
LHS				B	0.2772

Exhibit-1: For the Poisson Model, Hsu's MCB Results.

Hsu Individual 95% CIs For Mean Based on Pooled StDev						Hsu Intervals for Level Mean Minus Smallest of Other Level Means							
Level	N	Mean	StDev			Level	Lower	Center	Upper				
CMC	1920	0.3999	0.5749		(*-)	CMC	0.0000	0.3825	0.4083		(-----*)		
AV	1920	0.0795	0.1700		(*-)	AV	0.0000	0.0621	0.0879		(---*-)		
CRN	1920	0.0173	0.0411		(*-)	CRN	-0.0879	-0.0621	0.0000		(-*---)		
LHS	1920	0.1393	0.4863		(*-)	LHS	0.0000	0.1220	0.1477		(-----*-)		
				0.00	0.12	0.24	0.36			0.00	0.15	0.30	0.45

Exhibit-2: For the Negative Binomial Model, Hsu's MCB Results.

Hsu Individual 95% CIs For Mean Based on Pooled StDev						Hsu Intervals for Level Mean Minus Smallest of Other Level Means							
Level	N	Mean	StDev			Level	Lower	Center	Upper				
CMC	1920	0.998	1.984		(*-)	CMC	0.000	0.971	1.045		(-----*-)		
AV	1920	0.184	0.597		(*-)	AV	0.000	0.157	0.231		(---*--)		
CRN	1920	0.026	0.064		(*-)	CRN	-0.231	-0.157	0.000		(--*--)		
LHS	1920	0.287	0.779		(*-)	LHS	0.000	0.261	0.335		(-----*--)		
				0.00	0.30	0.60	0.90			0.00	0.35	0.70	1.05

Exhibit-3: For the Gamma Model, Hsu's MCB Results.

Hsu Individual 95% CIs For Mean Based on Pooled StDev						Hsu Intervals for Level Mean Minus Smallest of Other Level Means							
Level	N	Mean	StDev			Level	Lower	Center	Upper				
CMC	1920	1.141	2.390		(*-)	CMC	0.000	1.116	1.203		(-----*-)		
AV	1920	0.207	0.741		(*-)	AV	0.000	0.181	0.269		(----*-)		
CRN	1920	0.026	0.060		(*-)	CRN	-0.269	-0.181	0.000		(-*---)		
LHS	1920	0.277	0.851		(*-)	LHS	0.000	0.251	0.339		(-----*-)		
				0.00	0.35	0.70	1.05			0.00	0.40	0.80	1.20

In order to recommend the best VRT with a confidence, we apply the multiple comparison procedure so-called Hsu's multiple comparisons with the best (Hsu's MCB) (Hsu 1981). The method tests whether means are greater than the unknown minimum in case of determining the minimum. A statistically significant difference can only be observed between corresponding means if an interval contains zero as an end point. The results that are collected via the statistical package MINITAB by setting the default options are depicted in Exhibit 1, Exhibit 2 and Exhibit 3 under 95% confidence level for the Poisson, negative binomial and gamma models, respectively, with sample size 2000. As can be noted, CRN is the VRT whose performance is significantly different from others across all the demand models.

5 CONCLUSION AND FUTURE STUDY

This paper examines a constrained stochastic inventory optimization problem. For the solution of the problem, we describe an SAA based optimization procedure. In order to increase the efficiency of a potential algorithm for which the described optimization procedure is embedded, three variance reduction techniques; namely, antithetic variates, Latin hypercube sampling and common random numbers are evaluated for a set of test cases. We computationally investigate the impact of these VRTs on the optimization gap and gap variance by comparing the results with the crude Monte Carlo sampling method. The experiment results indicate that all three VRTs are affective at reducing the total gap variance, with the CRN outperforming other VRTs. One direction for future research is to investigate the applicability of other VRTs (e.g., importance sampling) for the underlying problem domain. Another direction lies in explicitly designing and testing the efficiency of the optimization algorithm within the proposed VRTs. The last, but not least, direction is to investigate the applicability of the procedure and variance reduction techniques for other inventory policies such as (s, S) and (r, nq) .

ACKNOWLEDGMENTS

We would like to thank Tito Homen-de-Mello for his helpful comments on this paper. This material is based upon work supported by the National Science Foundation under Grant No. 0437408. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Agrawal, V., and S. Seshadri. 2000. "Distribution free bounds for service constrained (r, q) inventory systems". *Naval Research Logistics* 47:635–656.
- Ahmed, S., and A. Shapiro. 2002. "The Sample Average Approximation Method for Stochastic Programs with Integer Recourse". Technical report, Georgia Institute of Technology.
- Axsäter, S. 2006. *Inventory Control*. Springer, New York.
- Federgruen, A., and Y. S. Zheng. 1992. "An Efficient Algorithm for Computing an Optimal (r, Q) Policy in Continuous Review Stochastic Inventory Systems". *Operations Research* 40:808–813.
- Freimer, M. B., D. J. Thomas, and T. Linderoth. 2010. "The impact of sampling methods on bias and variance in stochastic linear programs". *Computational Optimization and Applications* 46:1–25.
- Greenwald, A., B. Guillemette, V. Naroditskiy, and M. Tschantz. 2006. "Scaling Up the Sample Average Approximation Method for Stochastic Optimization with Applications to Trading Agents". *Lecture Notes in Computer Science* 3937:187–199.
- Hsu, J. C. 1981. "Simultaneous confidence intervals for all distances from the "best"". *Annals of Statistics* 9:1026–1034.

- Kleywegt, A. J., and A. Shapiro. 2001. "The sample average approximation method for stochastic discrete optimization". *SIAM Journal on Optimization*:502.
- Kramer, C. Y. 1956. "Extension of multiple range tests to group means with unequal number of replications". *Biometrics* 12:307 – 310.
- Law, A. M., and D. M. Kelton. 1999. *Simulation Modeling and Analysis*. 3rd ed. McGraw-Hill Higher Education.
- Linderoth, J., A. Shapiro, and S. Wright. 2006, February. "The empirical behavior of sampling methods for stochastic programming". *Annals of Operations Research* 142 (1): 215–241.
- Mak, W.-K., D. P. Morton, and R. K. Wood. 1999. "Monte Carlo bounding techniques for determining solution quality in stochastic programs". *Operations Research Letters* 24 (1-2): 47 – 56.
- Matala, A. 2008. "Sample Size Requirement for Monte Carlo – simulations using Latin Hypercube Sampling". Technical report, Helsinki University of Technology, Department of Engineering Physics and Mathematics.
- Rossetti, M., and Y. Ünlü. 2011. "Evaluating the Robustness of Lead Time Demand Models". *The International Journal of Production Economics* 134:159–176.
- Tukey, J. W. 1953. *The Collected Works of John W. Tukey VIII. Multiple Comparisons: 1948–1983*. Chapman & Hall, New York.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill, New York.

AUTHOR BIOGRAPHIES

YASİN ÜNLÜ is a Ph.D. candidate in the Industrial Engineering Department at the University of Arkansas. His dissertation research involves developing better inventory models for intermittent and highly variable demand in order to meet desired service level requirements. He has served as a research assistant on several logistics related projects funded by the National Science Foundation and the U. S. Air Force since 2006. He is an active member of IIE and INFORMS. His research interests include inventory control, network optimization and scheduling-sequencing. His email address is yunlu@uark.edu.

MANUEL D. ROSSETTI, Ph.D., P.E. is a Professor in the Industrial Engineering Department at the University of Arkansas. He received his Ph.D. in Industrial and Systems Engineering from The Ohio State University. He serves as an Associate Editor for the International Journal of Modeling and Simulation and is active in IIE, INFORMS, and ASEE. He was a WSC proceedings editor in 2004 and a co-editor for the WSC 2009 conference. He is also author of the textbook, *Simulation Modeling and Arena* published by John Wiley & Sons. His email address is rossetti@uark.edu.