ABSTRACT

Title of dissertation: THREE ESSAYS ON STOCHASTIC OPTIMIZATION APPLIED IN FINANCIAL ENGINEERING AND INVENTORY MANAGEMENT

Huiju Zhang, Doctor of Philosophy, 2007

Dissertation directed by: Professor Michael C. Fu
Robert H. Smith School of Business

Stochastic optimization methods are now being widely used in a multitude of applications. This dissertation includes three essays on applying stochastic optimization methods to solve problems in inventory management and financial engineering.

Essay one addresses the problem of simultaneous price determination and inventory management. Demand depends explicitly on the product price p, and the inventory control system operates under a periodic review (s, S) ordering policy. To minimize the long-run average loss, we derive sample path derivatives that can be used in a gradient-based algorithm for determining the optimal values of the three parameters (s, S, p) in a simulation-based optimization procedure. Numerical results for several optimization examples via different stochastic algorithms are presented, and consistency proofs for the estimators are provided.

Essay two considers the application of stochastic optimization methods to American-style option pricing. We apply a randomized optimization algorithm called Model Reference Adaptive Search (MRAS) to pricing American-style options through parameterizing the early exercise boundary. Numerical results are provided

for pricing American-style call and put options written on underlying assets following geometric Brownian motion and Merton jump-diffusion processes. We also price American-style Asian options written on underlying assets following geometric Brownian motion. The results from the MRAS algorithm are compared with the cross-entropy (CE) method, and MRAS is found to be an efficient method.

Essay three addresses the problem of finding the optimal importance sampling measure when simulating portfolios of credit risky assets. We apply a gradient-based stochastic approximation method to find the parameters in the minimum variance problem when importance sampling is used. The gradient estimator is obtained under the original measure. We also employ the CE method to solve the same variance minimization problem. Numerical results illustrating the variance reduction are presented for the estimation of the portfolios' expected loss, unexpected loss and quantiles.

THREE ESSAYS ON STOCHASTIC OPTIMIZATION APPLIED IN FINANCIAL ENGINEERING AND INVENTORY MANAGEMENT

By

Huiju Zhang

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2007

Advisory Committee:

Professor Michael C. Fu, Chair
Professor Bruce Golden
Professor Dilip Madan
Professor Paul Zantek
Professor Jianguo Liu

Dedicated to

My parents

Zhang Litao and Yang Hanfen

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor, Dr. Fu, for his invaluable instruction and advice throughout my graduate studies, especially for my dissertation work. He introduced me to the fascinating world of stochastic modeling. His enthusiasm and integral view on research has made a deep impression on me. Much of my work would have been impossible without his persistent guidance and encouragement.

I am deeply indebted to all the members of my dissertation committee – Dr. Golden, Dr. Liu, Dr. Madan, and Dr. Zantek – for taking effort in reading and providing a lot of insightful comments on earlier versions of this dissertation. I would also like to thank them for keeping an eye on the progress of my work and always being available when I needed their advice.

I owe an equally great debt to all the faculty members in my department for giving me training and knowledge. Thanks also to all my colleagues in the department for providing a good working atmosphere, and I benefited a lot from conversations with them.

Last but not least, I am very grateful to my parents for their understanding, patience and continuous support during the entire period of my study.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1

Introduction

Stochastic search and optimization techniques are widely used in a vast number of areas, including aerospace engineering, medicine, transportation, supply chain, statistics, and finance. Stochastic optimization refers to the minimization (or maximization) of a function in the presence of randomness in the optimization process and/or function evaluation. Stochastic methods are able to handle many problems for which deterministic optimization methods are inappropriate.

This dissertation applies stochastic optimization methods to solve problems in the field of financial engineering and inventory management. Specifically, the dissertation consists of three essays addressing the following problems: simultaneously determining optimal price and inventory levels in an (s, S) inventory system via stochastic approximation approaches; applying a model reference adaptive stochastic search algorithm to price American-style options; estimating the optimal measure change of importance sampling for portfolios of credit risky assets via a gradient-based stochastic approximation method and a random search algorithm.

1.1 Sample Path Derivatives for (s, S) Inventory Systems with Price Determination

This essay addresses an important problem in the interface between marketing and inventory planning, specifically that of simultaneously finding the optimal price and the optimal inventory control parameters in the face of uncertain price-dependent demands. In particular, we study a periodic-review, single-product inventory system with the objective of minimizing the average infinite-horizon loss rate (maximizing the average long-term profit rate), where the stationary demands faced by the system depend on the constant price p, and the system adopts an (s, S) control policy, in which an order is placed when and only when its inventory position falls below the level s, and the order amount is such that it will bring the inventory level up to S. Under certain conditions, an (s, S) policy has been proven to be optimal for inventory systems with a fixed ordering cost. Scarf (1960) showed that an (s, S) policy was optimal for the finite horizon dynamic inventory system in which the ordering cost was linear plus a fixed reorder cost and holding/penalty costs were convex. Clark and Scarf (1960) extended the results to multi-echelon inventory systems. Iglehart (1963) extended Scarf's study to the infinite horizon case and considered non-zero delivery lead-times, obtaining bounds on s and S, and investigating the limiting behavior of (s, S) pairs. Veinott and Wagner (1963) developed a computational approach for finding an optimal (s, S) inventory policy for the fully backlogged model with fixed set-up cost, linear purchase cost and i.i.d. discrete random demands. Hollier et al. (2005) applied algorithms based on branch-and-bound tree search technique and genetic algorithms to a modified (*s*, *S*) inventory system with lumpy demand items. A

detailed review on the evolution of inventory theory can be found in Scarf's (2002) paper.

Thomas (1974) incorporated pricing decisions into the (s, S) control policy, and the resultant strategy is referred to as an (s, S, p) policy. In this policy, the optimal price p is set to be contingent upon the inventory level and can change from one period to another. Federgruen and Heching (1999) characterized the structure of an optimal combined pricing and inventory strategy for both finite and infinite horizon models with variant price change restrictions. Feng and Xiao (2000) considered a continuous-time model with multiple prices and reversible changes in prices. They found that the optimal price level was based on the length of remaining sales time and on-hand inventory. More recently, Chen and Simchi-Levi (2002) used dynamic programming to determine price and inventory levels simultaneously at the beginning of each period. They showed that an (s, S, p) policy is optimal when the demand model is additive.

Dynamic pricing may not be desirable in some industries or for some companies. Under many circumstances, a more stable pricing policy than the aforementioned inventory-contingent ones is preferred, e.g., Wal-Mart's “Everyday Low Prices”. Also for mature products with stable demand that generally incorporate little seasonal effect or advanced technologies, there is relatively little price fluctuation, so the single price model is appropriate. Furthermore, Gallego and Van Ryzin (1994) showed that the optimal fixed-price policy is nearly as good as the optimal inventory-contingent one under rather mild assumptions. In this essay, we

assume there is a fixed price to be selected that influences the future demand levels in a known way, and an (s, S) policy is used to control the inventory.

According to Spall (2003), stochastic optimization algorithms are optimization algorithms that have one or both of the following properties: 1) There is random noise in the measurements of the criterion to be optimized and/or related information, such as the gradient vector of the criterion; 2) There is a random choice made in the search direction as the algorithm iterates toward a solution. The first property arises in simulation-based optimization where Monte Carlo simulations are run as estimates of an actual system. This class of algorithms includes Robbins-Monro (1951) stochastic approximation (RMSA), finite difference stochastic optimization (FDSA), and simultaneous perturbation stochastic optimization (SPSA) introduced by Spall (1992). Algorithms satisfying the second property include genetic algorithms, simulated annealing, and random search algorithm (Zhigljavsky, 1991). Reviews of techniques for simulation-based optimization can be found in Jacobson and Schruben (1989), Safizadeh (1990), and Fu (1994, 2002); see also Spall (1999) for a detailed review on stochastic optimization.

Three stochastic optimization algorithms are used to find the optimal parameters in this essay: RMSA, a gradient-based search algorithm, SPSA, and simulated annealing (Kirkpatrick et al., 1983). To use gradient-based optimization requires sample path derivatives, the main focus of our work.

Fu (1994a) developed sample path derivatives using perturbation analysis (PA) for an inventory system adopting the (s, S) control policy, and Fu and Healy (1997) investigated their use in simulation-based optimization. Systematic and thorough

reviews on gradient estimation via perturbation analysis can be found in Glasserman (1991), Ho and Cao (1991), Fu and Hu (1997), and Fu (2006). We also use PA to derive our sample path derivatives. Similar to Fu's (1994a) approach, we implement infinitesimal perturbation analysis (IPA) and smoothed perturbation analysis (SPA) to estimate the derivatives. In Fu's model, demand is assumed to be exogenously specified, whereas in our model it depends on product price, which allows demand to be adjusted according to product properties such as price elasticity of demand. The inclusion of price in the model makes the derivation more complicated.

1.2 Applying Model Reference Adaptive Search to American-style Option Pricing

Pricing American option is a challenging problem in the financial engineering due to the early exercise features. Because of the complexity of the underlying dynamics, analytical models for option pricing entail many restrictive assumptions. Indeed, there is no analytical solution for the valuation of an American option on a single dividend-paying asset in the standard Black-Scholes framework. Methods other than applying simulation include lattice methods such as binomial and trinomial trees, and finite difference methods to solve the associated boundary condition partial differential equations (PDEs). In general, the computational speed of these methods is significantly better than that of simulation methods for simple models; however, these methods often only handle limited number of uncertainty sources in low dimension and become impractical in situations where there are multiple factors. For instance, pricing an Asian option is generally required to solve a PDE in two space dimensions,

which is prone to oscillatory solutions. In contrast, Monte Carlo simulation methods are more widely applicable, because they have no problem in dealing with high dimension and can manage complicated derivatives with more state variables. A number of simulation-based approaches have been developed to price American options since the 1990s.

A standard Monte Carlo simulation generates final payouts independently and averages them to obtain the expected value, while early exercise requires knowledge of the option value at intermediate dates and performs non-linear operations along the way. We classify Monte Carlo simulation algorithms designed to handle early exercise features into three main categories. The first class casts the problem in a stochastic dynamic programming framework and employs a backwards induction algorithm. At each early exercise date, the payoff from immediate exercise is compared to the holding value, i.e., the conditional expectation from keeping the derivative alive. However, computing this conditional expectation can become computationally prohibitive as the dimension of the problem increases, and the next-stage value function is calculated over its entire asset space domain. Tilley (1993) first applied a bundling technique to approximate the holding values at early exercise points. Improvements on the Tilley's methods include Carriere (1996), who used a spline and local regression technique to approximate the conditional expectations and find the optimal stopping in finite discrete time; and Longstaff and Schwartz (2001), who used least-square regression to provide a direct estimate of the conditional expectation function in high-dimensional setting. In addition, Tsitsiklis and Van Roy (2001) provided theoretical results that help explain the success of approximate

dynamic programming methods. Laprise et al. (2006) applied secant and tangent interpolations to construct a piecewise linear approximation of the value function, and estimate the American-style derivative by pricing a portfolio of European options at varying strike prices.

The second class of algorithms characterizes the optimal early exercise policies directly. Grant et al. (1996, 1997) identified the optimal critical price, i.e., the price below (above) which it is optimal to exercise for American put (call), using the backward recursive technique of dynamic programming and incorporated this early exercise feature into Monte Carlo simulation. Fu and Hu (1995) cast the American option pricing problem as an optimization problem of maximizing the expected payoff with respect to the early exercise thresholds, and incorporated the gradient estimates in an iterative stochastic approximation algorithm. Fu et al. (2001) presented another way to solve this optimization problem using simultaneous perturbation stochastic approximation (SPSA) proposed by Spall (1992).

The third class of algorithms is based on obtaining upper and lower bounds from simulated paths and backwards recursion. Broadie and Glasserman (1997) proposed a method based on simulated nonrecombining trees, where both bounds converge to the true price as computational effort increases. Broadie and Glasserman (2004) introduced a stochastic mesh method for pricing high-dimensional American options with a finite, but possibly large, number of exercise dates. The computational effort of this algorithm is linear in its dependence on the number of exercise dates, in contrast to the exponential dependence for random tree method.

Other versions of approximate value iteration have also been proposed in the options pricing literature. Some involve partitioning the state space and computing one value per partition. Barraquand and Martineau (1995) combined Monte Carlo simulation with stratified state aggregation techniques to approximate the price of American securities. This can be viewed as a version of approximate value iteration involving piecewise constant approximations. Furthermore, Keber (1999) implemented a genetic programming approach to derive a formula for American put options and showed that genetically determined formulas outperformed most frequently quoted analytical approximations in calculating the implied volatility based on the Black-Scholes model.

In this essay, we apply a randomized algorithm called Model Reference Adaptive Search (MRAS) for pricing American-style options by solving an optimization problem in the spirit of the second class of algorithms discussed above. We compare our numerical results with those computed from perturbation analysis stochastic approximation (PASA) and SPSA approaches described in Fu et al. (2001).

MRAS was proposed by Hu et al. (2007). The main idea of this approach is similar to that of the cross-entropy (CE) method (Rubinstein and Kroese 2004), which has been successfully applied to a wide range of combinatorial optimization and rare-event estimation problems. In contrast to instance-based methods such as simulated annealing (Aarts and Korst, 1989), threshold acceptance (Dueck and Scheur, 1990), genetic algorithms (GAs) (Srinivas and Patnaik 1994) and tabu search (Glover 1990), where the new candidate solutions generated in the next iteration depend directly on solution or the 'population' of solutions from previous step, both MRAS and CE fall

in the category of model-based search algorithms, which construct a random sequence of solutions via an intermediate probabilistic model that is updated from the previous solutions in such a way that the search will concentrate in the regions containing high quality solutions, and usually involve the following two iterative phases:

1. Generate candidate solutions (random data samples, vectors, trajectories, etc.) according to a specified random mechanism, e.g., a parameterized probability distribution.
2. Update the parameters of the random mechanism, typically parameters of pdfs, on the basis of the data collected in the previous step, to produce a "better" sample of candidate solutions in the next iteration.

The obtained parameters tend to coincide with the parameters that minimize variance in most cases such that the outcome converges probabilistically to the optimal or near-optimal solution.

1.3 Optimizing Importance Sampling Parameter for Portfolios of Credit Risky Assets

Credit risk modeling has gained increasing interest among bankers and other portfolio managers since the mid-1990s, and the development of market risk management measures such as value-at-risk (VAR) has accelerated this approach. Accurate assessments of the risk of large potential losses on a credit portfolio play a key role in the management of financial institutions with large credit portfolios. Credit risk is the risk due to uncertainty in counterparty's ability to meet its obligations. The counterparty could be an individual, a corporation or financial institution, or a

sovereign government. The most common corporate credit instruments are bonds and loans.

Two primary types of credit risk have been described in the literature. The first — structural credit risk models — identify the loss in the portfolio if the obligor has defaulted on its legal obligations within a certain time horizon. The first structural model was proposed in Black and Scholes' (1973) influential paper on option pricing and was studied in more detail in Merton (1974). In Merton's model, default probabilities are calculated on the basis of a firm's capital structure and asset volatility. A firm defaults when the value of its liabilities exceeds its assets at the debt's maturity date, and it uses an option framework to calculate this risk neutral default probability. A popular implementation of this model is the commercial KMV (Kealhofer, McQuown and Vasicek) model that is the foundation of our study in this essay. Another approach, within the structural framework, was introduced by Black and Cox (1976). In contrast to the Merton model, defaults occur as soon as the firm's asset value falls below a certain threshold in this model, that is, default can occur at any time. The structural models allow default hedging, but are difficult to calibrate to the market data.

The second method — reduced form models, also called intensity models — is marked to market. They recognize any gains or losses in the value of a debt security caused by changes in the credit quality of the obligor over the measured time horizon. Reduced form models use market prices of the firms' defaultable instruments such as bonds or credit default swaps, to extract both their default probabilities and their credit risk dependencies. In contrast to the structural models, they rely on the

market as the only source of information without considering any information included in balance sheets or equity prices. A portfolio's market value resulting from defaults or changes in credit ratings are modeled, and the time of default is not determined via the asset and liability value of the firm, but by an exogenously given jump process. The first published intensity model appears to be Jarrow and Turnbull (1995). Subsequent research includes Duffie and Huang (1996), Jarrow et al. (1997), and Duffie and Singleton (1997). Madan and Unal (1998) presented one of the first intensity-based credit risk models. In their model the event of default has two underlying risk components, one associated with the timing of the event and the other with its magnitude. The default intensity is directly linked to the market value of the firm's equity.

Despite the different methodologies, all credit risk models develop a distribution of possible credit portfolio values at some point in the future. Correlated changes in the credit quality of underlying risky assets result in changes in the value of exposures. These exposures are then aggregated to produce the portfolio loss distribution, which indicates the probability of achieving a certain portfolio value over a certain time period. Most of the time the risky asset does not default and the loss is zero; however, the loss is usually substantial when default occurs. Consequently, the distribution of possible future losses for a portfolio of credit risky assets shows strongly asymmetric behavior and a fat tail as a consequence of the limited upside of credit (the promised coupon payment) and substantial but rare downside if the corporation defaults. It is not possible to fully diversify away the fat tail because of correlation. There is always a large probability of relatively small

losses and a small probability of rather large losses. Because of the complexity of the portfolio structure, Monte Carlo simulation is widely applied to determine the loss distribution for a credit portfolio. However, it is usually time consuming to apply Monte Carlo simulation to provide sufficiently accurate estimation. The size of the portfolios and the complexity of the assumptions make speed issues particularly acute. Therefore, variance reduction methods that can speed up the computation are of substantial interest.

Importance sampling (IS) method is a common variance reduction technique for increasing the efficiency of Monte Carlo simulation. The basic idea is to focus simulation effort on the most important regions of the space from which samples are drawn. The simulation outputs are weighted to correct for the use of the biased distribution, and this ensures that the new IS estimator is unbiased. The weight is given by the likelihood ratio, that is, the Radon-Nikodym derivative of the true underlying distribution with respect to the biased simulation distribution. An overview of importance sampling methodology can be found in Glasserman (2003). Some recent work on applying IS in credit risk portfolios include Kalkbrener et al. (2003), Morokoff (2004), and Glasserman et al. (2005).

The fundamental issue in implementing IS simulation is the choice of the biased distribution that enhances the frequency of the important events, and compensates the bias through the multiplication by the likelihood ratio. The rewards for a good distribution can be huge run-time savings; the penalty for a bad distribution can be longer run times than for a standard Monte Carlo simulation without any special techniques. Vazquez-Abad and Dufresne (1998) employed

gradient-based method to estimate the optimal importance sampling measure in pricing Asian options. Su and Fu (2002) proposed a similar stochastic approximation method but using a more general gradient estimator derived under a different measure. Because we can cast choosing the optimal measure change of importance sampling as an optimization problem, we will use stochastic search algorithms such as CE method to find the optimal solution, as well. More studies on the selection of an importance sampling change of measure that leads to an efficient variance reduction will be presented in this essay.

1.4 Research Contributions

Stochastic optimization methods have been playing a rapidly grow role in the analysis, design, and operation of modern systems. They provide a way of handling inherent system noise and models that are inappropriate for classical deterministic methods of optimization. In this dissertation, we studied extensively the application of stochastic optimization in the field of inventory management and financial engineering. The main contributions are:

i) We apply a stochastic approximation algorithm to find optimal price and inventory levels for an (s, S) system with price determination. The gradient estimators are derived and used in an iterative gradient-based optimization algorithm. Consistency proofs for the estimators are provided. We apply three stochastic optimization algorithms to the (s, S, p) inventory model and study the behavior of optimal parameters. The results

from the gradient-based algorithm RMSA are compared with those from SPSA and simulated annealing approaches. We study the effect of unit holding cost, fixed ordering cost and price elasticity of demand in estimating optimal inventory levels and selling price. Price is found to be the determining factor, because revenue is much larger than the fixed ordering cost and holding costs in all the numerical cases. We also demonstrate that RMSA is most efficient in terms of convergence performance, because of the gradient information.

We consider a single product model in our studies, while the realistic problems are more complicated with possible correlations between similar products or dynamic pricing involved. There are numerous software packages available for merchandise optimization in industry, e.g., the leading retail software vendor ProfitLogic provides Retail Profit Optimization solutions to manage retailers' pricing and discounting policies in order to get the most revenue possible on their inventories. Compared with those real-world models, only a simple case with three parameters is studied in this essay; however, our model provides a framework to determine the inventory levels and fixed price simultaneously under some mild conditions and it is extendable.

ii) We apply a stochastic optimization method called Model Reference Adaptive Search to price American-style options through parameterizing the early exercise thresholds. The optimal values are reached for various

cases, including American call option, American put option, and path dependent American-style Asian options, assuming the underlying asset follows geometric Brownian motion or the Merton jump-diffusion process. We demonstrate that the global maximum is consistently found for varying initial condition settings and conclude that MRAS can be a very effective approach. The optimization algorithm presented in this essay is widely applicable, and not just limited to the cases we have studied.

iii) We provide a general framework to find optimal measure change in employing importance sampling technique in Monte Carlo simulation. To reduce variance in estimating loss of a portfolio of credit risky assets by Monte Carlo simulation, we formulate a parametric minimization problem for the optimal importance sampling measure. We apply a gradient-based stochastic approximation algorithm via infinitesimal perturbation analysis and the CE method to solve the optimization problem. Both algorithms converge efficiently to the optimum and yield a significant improvement in accuracy in estimating the expected loss, unexpected loss, and quantiles. We also show that the objective function is a convex function. The stochastic optimization approaches we present in this essay are capable of finding the optimal change of importance sampling measure efficiently. We found significant variance reductions from the simulation results, with the computational gains ranging from 2.0 to 82.6 in our numerical examples.

Chapter 2

Sample Path Derivatives for (s, S) Inventory Systems with Price Determination

This essay is organized in the following manner. Section 1 reviews the (s, S, p) model and the demand structure. The IPA analysis is presented in Section 2, and the SPA analysis is developed in Section 3. Section 4 contains the consistency proof for the infinite horizon model. Section 5 presents a numerical example where the estimators are used in a gradient-based algorithm to search for the optimal setting of the parameters (s, S, p). The results are compared with those from SPSA and simulated annealing. Section 6 concludes the essay with a brief summary.

2.1 Model Formulation

Consider a firm that has to make production and price decisions under stationary independently and identically distributed (i.i.d.) demand that depends on a constant product price. For each period t, t = 1,2 …, T, let

D_t := demand in period t, i.i.d., with p.d.f. f and c.d.f. F,

p: = selling price.

where the demand function is of the general form

$$D_t = d(p, \varepsilon_t) := \varepsilon_t \gamma(p) + \delta(p) \tag{2.1}$$

with $\gamma(.)$ and $\delta(.)$ nonincreasing functions and ε_t assumed to be i.i.d. The cases $\gamma(p) = 1$ and $\delta(p) = 0$ are often referred to as the additive and multiplicative model, respectively. We use additive stochastic demand functions in our model with $\delta(p) = b - a*p$, a, b > 0, where a is the price elasticity of demand, which measures the nature and degree of consumers' respond in their buying decisions to a change in product price.

In this essay, we assume that the ordering decision is made at the beginning of the period, and the demand for the period is subtracted at the end of the period. Let x_t be the inventory level at the beginning of period t before placing an order, and y_t be the inventory level at the beginning of period t after placing an order. Hence, $y_t = x_t$ if no order is made, and $y_t = S > x_t$ if order is placed at beginning of period t. The ordering cost includes both a fixed cost and a variable cost proportional to the amount ordered. Demand that cannot be met from inventory on hand is fully backordered. The inventory carrying and stockout costs all depend on the size of the end-of-the-period inventory level and shortfall. In addition, we assume the order lead time is zero

throughout this essay, i.e., no delay between the placing of an order and the receipt of the goods ordered, thus, the inventory position coincides with the inventory level. We can easily obtain our derivatives with a fixed lead time by rewriting the expression between the inventory position and inventory level, as shown in Fu (1994a). The objective of this study is to find underlying parameters θ = (s, S, p) to minimize the long-run average loss L:

$$\begin{aligned} L_T &= \frac{1}{T}\sum_{t=1}^{T} l(y_t, p) \\ &= \frac{1}{T}\sum_{t=1}^{T} [I\{x_t < s\}(k + c(S - x_t)) + h(y_t - D_t)^+ + g(y_t - D_t)^- - pD_t(p, \varepsilon_t)] \end{aligned} \tag{2.2a}$$

$$L = \lim_{T\to\infty} L_T \tag{2.2b}$$

where k is the fixed cost of placing an order, c is the variable order cost coefficient, h is the holding cost coefficient, g is the shortage cost coefficient, p is the revenue coefficient, I{.} is the indicator function, and $x^+ = \max(0, x)$, and $x^- = \max(0, -x)$.

2.2 IPA Estimation

Our goal in this section is to develop derivative estimators for L_T with respect to the control parameters θ, where θ = s, S, and p. Without loss of generality, we assume that $y_0 = S$. We define q = S - s for notational convenience in the analysis that follows.

According to the definition of x_t and y_t, the recursive dynamic equation for y_t is given by

$$y_t = \begin{cases} x_t = y_{t-1} - D_{t-1} & \text{if } x_t \geq s \\ S & \text{if } x_t < s \end{cases} \tag{2.3}$$

That is to say, $y_t = x_t$ if the beginning inventory level is greater than the reorder point s; otherwise, an order is placed so that $y_t = S$. Thus

$$\frac{\partial y_t}{\partial \theta} = \begin{cases} \frac{\partial y_{t-1}}{\partial \theta} - \frac{\partial D_{t-1}}{\partial \theta} & if\ x_t \geq s \\ \frac{\partial S}{\partial \theta} & if\ x_t < s \end{cases} \tag{2.4}$$

With the initial condition $y_0 = S = s + q$, we have $\partial y_0 / \partial s = \partial y_0 / \partial q = 1, \partial y_0 / \partial p = 0$. According to the equation (2.4), we have $\partial y_t / \partial s = \partial y_t / \partial q = 1$ for all t, because $D_t = b - a * p + \varepsilon_t$, which is independent of s or q. Thus,

$$\frac{\partial y_t}{\partial p} = \begin{cases} \frac{\partial y_{t-1}}{\partial p} + a & if\ x_t \geq s \\ 0 & if\ x_t < s \end{cases} \tag{2.5}$$

By applying above recursive dynamic equation backwards, for the no-order-decision-made-period, if $x_t \geq s,, x_{t-1} \geq s, ... x_{t'+1} \geq s, x_{t'} < s$, we have

$$\frac{\partial y_t}{\partial p} = \frac{\partial y_{t-1}}{\partial p} + a = (\frac{\partial y_{t-2}}{\partial p} + a) + a = ... = \frac{\partial y_{t'}}{\partial p} + (t - t')a = (t - t')a,$$

where t' is the most recent period that an order is placed before t. So we can rewrite equation (2.5) as

$$\frac{\partial y_t}{\partial p} = \begin{cases} (t - t')a & x_t \geq s, x_{t-1} \geq s ... x_{t'+1} \geq s, x_{t'} < s \\ 0 & x_t < s \end{cases} \tag{2.6}$$

If we place an order in period t, the inventory level will be brought back to S in the same period since ordering lead time is zero. Recall that ordering cost consists

of a fixed set-up cost and a variable cost proportional to the ordering amount, so the ordering cost in period t can then be formulated as

$$k\delta(y_t - x_t) + c(y_t - x_t) = \begin{cases} 0 & x_t \geq s \\ k + c(S - x_t) & x_t < s \end{cases} \tag{2.7}$$

The first derivative of ordering cost in period t is straightforward:

$$\frac{\partial(k\delta(y_t - x_t) + c(y_t - x_t))}{\partial\theta} = \begin{cases} 0 & x_t \geq s \\ c\dfrac{\partial(S - x_t)}{\partial\theta} & x_t < s \end{cases} \tag{2.8}$$

When applying the recursive dynamic relation described in equation (2.4) and (2.5) for the case of $x_t < s$, we have

$$\begin{aligned} & c\frac{\partial(S - x_t)}{\partial\theta} = c\frac{\partial(S - (y_{t-1} - D_{t-1}))}{\partial\theta} = c(\frac{\partial S}{\partial\theta} - \frac{\partial y_{t-1}}{\partial\theta} + \frac{\partial D_{t-1}}{\partial\theta}) \\ & = \begin{cases} c(-\dfrac{\partial y_{t-1}}{\partial\theta} - a) = -c(t - t')a & if \quad \theta = p \\ c(1 - \dfrac{\partial y_{t-1}}{\partial\theta}) = 0 & if \quad \theta = s, q \end{cases} \end{aligned} \tag{2.9}$$

The derivative of holding cost with respect to s and q is the holding cost coefficient itself for all the periods, considering the demand is independent on s or q, and $\frac{\partial y_t}{\partial s} = \frac{\partial y_t}{\partial q} = 1$, specifically, $\frac{h\partial(y_t - D_t)}{\partial\theta} = h\frac{\partial y_t}{\partial\theta} = h$. Similarly, the derivative of shortage cost to s and q is -g. Applying equation (2.6), we obtain the direct differentiation of holding/shortage cost to price p in period t:

$$\frac{\partial h(y_t - D_t)}{\partial p} = h(\frac{\partial y_t}{\partial p} - \frac{\partial D_t}{\partial p}) = \begin{cases} h(t - t'+1)a & x_t \geq s \\ ha & x_t < s \end{cases} \tag{2.10}$$

And

$$\frac{\partial g(D_t - y_t)}{\partial p} = g(\frac{\partial D_t}{\partial p} - \frac{\partial y_t}{\partial p}) = \begin{cases} -g(t-t'+1)a & x_t \geq s \\ -ga & x_t < s \end{cases} \tag{2.11}$$

The sample path derivative of the revenue term $p*D_t$ in period t with respect to s and q is 0, since they rely only on price, not stock levels, whereas the sample path derivatives with respect to p is $-2ap+b+\varepsilon_t$. Combining all the analyses, we obtain the complete IPA estimator for the time-average loss:

$$(\frac{\partial L_T}{\partial \theta})_{IPA} = \frac{1}{T}\sum_{t=1}^{T}\frac{\partial l(y_t, p)}{\partial \theta} = \frac{1}{T}(\sum_{t:x_{t+1}>0} h - \sum_{t:x_{t+1}<0} g) \qquad for \quad \theta = s, q, \tag{2.12}$$

$$\begin{aligned}(\frac{\partial L_T}{\partial p})_{IPA} &= \frac{1}{T}\sum_{t=1}^{T}\frac{\partial l(y_t, p)}{\partial p} \\ &= \frac{1}{T}[\sum_{t:x_t<s} -c(t-t')a + \sum_{t:x_t\geq s, x_{t+1}>0} h(t-t')a + \sum_{t:x_{t+1}>0} ha - \sum_{t:x_t\geq s, x_{t+1}<0} g(t-t')a \\ &\quad - \sum_{t:x_{t+1}<0} ga - \sum_{t=1}^{T}(-2ap+b+\varepsilon_t)]\end{aligned} \tag{2.13}$$

2.3 SPA Estimation

In sample path analysis of the (s, S) model, Fu (1994a) concluded that IPA alone is sufficient for estimating the derivative with respect to θ = s, but not for θ = q, where an additional SPA (smoothed perturbation analysis) term must be added. This conclusion also holds for our model, so we need to use SPA.

We consider a positive change in s. Fig.2.1 shows the perturbation path for a small positive change Δs in the reorder point s. The sample path moves upward by Δs smoothly, i.e., the sample performance is continuous, so IPA alone suffices for s

(assuming q is held constant). However, for $\Delta q > 0$, it is possible that an ordering decision changes from order to not order in a period. Fig.2.2 represents the sample path for change in q, and period t is the order-decision-change period. Since Δq is an infinitesimal amount and demand is finite, the demand during t will lead to an order decision in the next period. The perturbed path for inventory position can be constructed from the nominal path with an appropriate extra period "inserted"'. The beginning inventory in this period is $y = s - \alpha + \Delta q$. Then, an SPA term based upon conditional expectation is added to smooth the discontinuity:

$$\begin{aligned}(\frac{\partial L_T}{\partial q})_{SPA} &= \sum_{t\in M^*(T)} \lim_{\Delta q\to 0} \frac{E_{z_t}[\Delta L_T | \alpha_t \le \Delta q] P_{z_t}\{\alpha_t \le \Delta q\}}{\Delta q} \\ &= \sum_{t\in M^*(T)} \lim_{\Delta q\to 0} E_{z_t}[\Delta L_T | \alpha_t \le \Delta q] \bullet \lim_{\Delta q\to 0} \frac{P_{z_t}\{\alpha_t \le \Delta q\}}{\Delta q}\end{aligned} \tag{2.14}$$

where $Z_t = y_{t-1} - s$, $\alpha_t = D_{t-1} - Z_t$, and $M^*(T) = \{t \le T : y_t = S\}$ is the set of periods in which orders are placed.

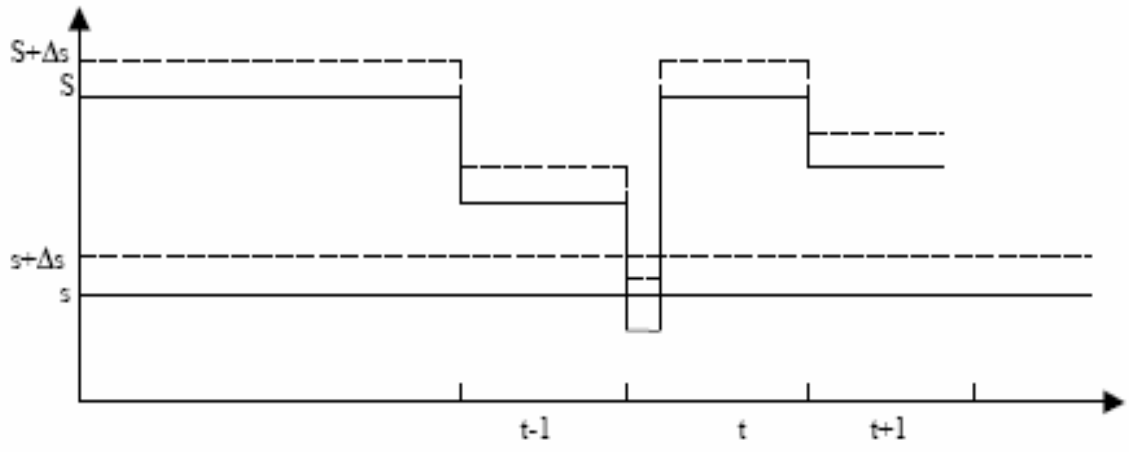

Figure 2.1. Effect on sample path with p, q fixed and s perturbed

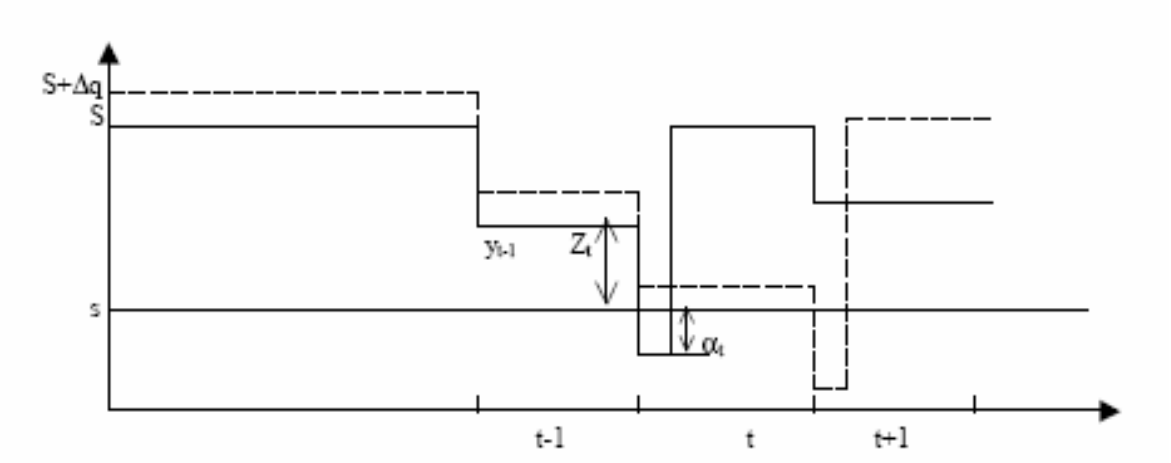

Figure 2.2. Effect on sample path with s, p fixed and q perturbed

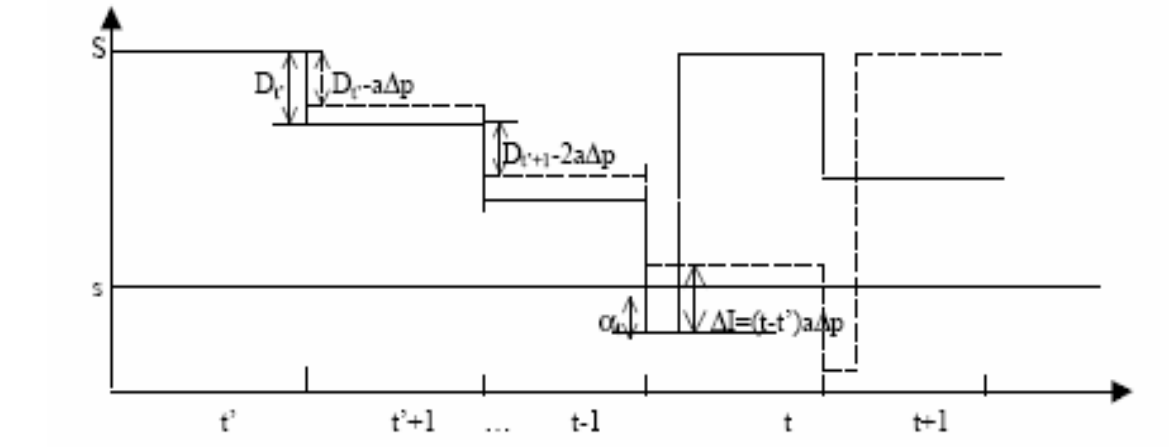

Figure 2.3. Effect on sample path with s, q fixed and p perturbed

Note that in the rest of the derivation, we will often drop the subscripts for notational convenience. The latter term can be estimated explicitly from the original sample path, given the demand distribution:

$$P_z(\alpha \le x \mid D > z\} = P_z\{D - z \le x \mid D > z\} = \frac{P_z\{z < D \le z + x\}}{P_z\{D > z\}} = \frac{F(z+x) - F(z)}{1 - F(z)}.$$

Hence,

$$\lim_{\Delta q \to 0} \frac{P_z\{\alpha \le \Delta q\}}{\Delta q} = \frac{f(z)}{1-F(z)} \tag{2.15}$$

We have the performance measure from the perturbed path as

$$E_z(L_{T+1}(q+\Delta q) \mid \alpha \le \Delta q) = \frac{1}{T+1}[\sum_{t=1}^{T} l(y_t + \Delta q, p) + E[l(s-\alpha+\Delta q, p) \mid \alpha \le \Delta q]]$$

$$= \frac{1}{T+1}[\sum_{t=1}^{T} l(y_t + \Delta q, p) + cE[D] + hE[(s-D-\alpha+\Delta q)^+] + gE[(s-D-\alpha+\Delta q)^-] - pE[D]],$$

and

$$E_z[\Delta L_T \mid \alpha \le \Delta q] = E_z[L_{T+1}(q+\Delta q) - L_T(q) \mid \alpha \le \Delta q]$$

$$= \frac{1}{T+1}\{\sum_{t=1}^{T}[l(y_t + \Delta q, p) - l(y_t, p)] + cE[D] + hE[(s-D-\alpha+\Delta q)^+]$$

$$+ gE[(s-D-\alpha+\Delta q)^-] - pE[D] - \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\}$$

Take $\Delta q \to 0$, we get

$$\lim_{\Delta q \to 0} E_z[\Delta L_T \mid \alpha \le \Delta q]$$

$$= \frac{1}{T+1}\{cE[D] + hE[(s-D)^+] + gE[(s-D)^-] - pE[D] - \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\}. \tag{2.16}$$

Incorporating (2.15) and (2.16) and reinstalling our subscripts into (2.14), we have

$$(\frac{\partial L_T}{\partial q})_{SPA} = \frac{1}{T+1}\{cE[D] + hE[(s-D)^+]$$

$$+ gE[(s-D)^-] - pE[D] - \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\} \bullet \sum_{t \in M^*(T)} \frac{f(z_t)}{1-F(z_t)}.$$

The final estimator for the derivative with respect to q is the sum of the IPA and SPA parts:

$$(\frac{\partial L_T}{\partial q})_{PA} = \frac{1}{T}(\sum_{t:x_{t+1}>0} h - \sum_{t:x_{t+1}<0} g) + \frac{1}{T+1}\{cE[D] + hE[(s-D)^+] \\ + gE[(s-D)^-] - pE[D] - \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\} \bullet \sum_{t\in M^*(T)} \frac{f(z_t)}{1-F(z_t)}. \tag{2.17}$$

As the price changes from p to p+Δp, demand will decrease in each period by the amount of aΔp. Figure 2.3 illustrates the sample path with Δp>0. An additional SPA term is needed for estimator with respect to p, since the order decision may change. The sample path is similar to that for q, but instead of a change Δq, there is an accumulated change $\Delta I = (t-t')a\Delta p$, where $t \in M^*(T)$. The derivative analysis is similar to the analysis we had done for the perturbation path with a change of Δq. First we define $\beta_t(p) = D_{t'}(p) + ... + D_{t-1}(p) - q = \varepsilon_{t'} + ... + \varepsilon_{t-1} + (t-t')(b-ap) - q$. Then we have

$$(\frac{\partial L_T}{\partial p})_{SPA} = \sum_{t\in M^*(T)} \lim_{\Delta p\to 0} \frac{E_{z_t}[\Delta L_T | \beta_t(p) \le \Delta I] P_{z_t}\{\beta_t(p) \le \Delta I\}}{\Delta p} \\ = \sum_{t\in M^*(T)} \lim_{\Delta p\to 0} E_{z_t}[\Delta L_T | \beta_t(p) \le \Delta I] \bullet \lim_{\Delta p\to 0} \frac{P_{z_t}\{\beta_t(p) \le \Delta I\}}{\Delta p} \tag{2.18}$$

Let $f_\varepsilon(x) = f(x+b-ap)$ and $F_\varepsilon(x) = F(x+b-ap)$ the p.d.f. and c.d.f., respectively, of ε_t. We have

$$P(\beta_t(p) \le x \mid D_{t'}(p) + ... + D_{t-2}(p) = q - z_t, D_{t-1}(p) > z_t) \\ = P(\varepsilon_{t-1} \le z_t - b + ap + x \mid \varepsilon_{t'} + ... + \varepsilon_{t-2} = q - z_t + (t-t'-1)(-b+ap), \varepsilon_{t-1} > z_t - b + ap) \\ = \frac{F_\varepsilon(z_t - b + ap + x) - F_\varepsilon(z_t - b + ap)}{1 - F_\varepsilon(z_t - b + ap)} = \frac{F(z_t + x) - F(z_t)}{1 - F(z_t)},$$

where the second equality is due to the independence between ε_t's. Therefore,

$$\lim_{\Delta p\to 0} \frac{P_{z_t}\{\beta_t(p) \le \Delta I\}}{\Delta p} = \frac{f(z_t)}{1-F(z_t)} \bullet (t-t')a.$$

We also notice that the inserted period has inventory level y = s – β + ΔI, so

$$
\begin{aligned}
&E_z(L_{T+1}(p+\Delta p)\,|\,\beta \le \Delta I) \\
&= \frac{1}{T+1}[\sum_{t=1}^{T} l(y_t, p+\Delta p) + E[l(s-\beta+\Delta I, p+\Delta p)\big|\beta \le \Delta I]] \\
&= \frac{1}{T+1}[\sum_{t=1}^{T} l(y_t, p+\Delta p) + cE[D] + hE[(s-D-\beta+\Delta I)^+] + gE[(s-D-\beta+\Delta I)^-] - (p+\Delta p)E[D]],
\end{aligned}
$$

and

$$
\begin{aligned}
&E_z(\Delta L_T)\,|\,\beta \le \Delta I) \\
&= \frac{1}{T+1}\{\sum_{i=1}^{T}(l(y_t, p+\Delta p) - l(y_t, p)) + cE[D] + hE[(s-D-\beta+\Delta I)^+] \\
&+ gE[(s-D-\beta+\Delta I)^-] - (p+\Delta p)E[D] - \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\},
\end{aligned}
$$

$$
\begin{aligned}
&\lim_{\Delta p \to 0} E_z[\Delta L_T\,|\,\beta \le \Delta I] \\
&= \frac{1}{T+1}\{cE[D] + hE[(s-D)^+] + gE[(s-D)^-] - pE[D] - \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\},
\end{aligned}
\tag{2.19}
$$

Therefore,

$$
\begin{aligned}
(\frac{\partial L_T}{\partial p})_{SPA} = &\frac{1}{T+1}\{cE[D] + hE[(s-D)^+] + gE[(s-D)^-] - pE[D] \\
&- \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\} \sum_{t \in M^*(T)} [\frac{f(z_t)}{1-F(z_t)}(t-t')a]
\end{aligned}
\tag{2.20}
$$

Summation of IPA and SPA terms gives derivate of loss function to p:

$$
\begin{aligned}
&(\frac{\partial L_T}{\partial p})_{PA} = \frac{1}{T}[\sum_{t:x_t<s} -c(t-t')a + \sum_{t:x_t \ge s, x_{t+1}>0} h(t-t')a + \sum_{t:x_{t+1>0}} ha - \sum_{t:x_t \ge s, x_{t+1}<0} g(t-t')a \\
&- \sum_{t:x_{t+1<0}} ga - \sum_{t=1}^{T}(-2ap+b+\varepsilon_t)] + \frac{1}{T+1}\{cE[D] + hE[(s-D)^+] + gE[(s-D)^-] - pE[D] \\
&- \frac{1}{T}\sum_{i=1}^{T} l(y_t, p)\} \sum_{t \in M^*(T)} [\frac{f(z_t)}{1-F(z_t)}(t-t')a]
\end{aligned}
\tag{2.21}
$$

2.4 Consistency Proof

Our goal is to prove the consistency:

$$(\frac{\partial L_T}{\partial \theta})_{PA} \to \frac{\partial E[l(y,p)]}{\partial \theta},$$

where $l(y_t, p) = (k + c(S - x_t)) \cdot \delta(y_t - x_t) + h(x_{t+1})^+ + g(x_{t+1})^- - pD_t$, loss in the period t.

Let X and Y denote the steady-state random variables for x_t and y_t. Then, we have

$\sum_{t-1}^{T} x_t / T \to E[X]$, $\sum_{t-1}^{T} y_t / T \to E[Y]$, and $\sum_{t-1}^{T} D_t / T \to E[D]$. From Fu's derivation (1994a), we already know

$$P(Y = S) = 1/(1 + R(q)),$$

and

$$E[Y] = s + \frac{q + \overline{R}(q)}{1 + R(q)}.$$

In our model, we have X=Y-D, therefore

$$E[X] = s + \frac{q + \overline{R}(q)}{1 + R(q)} - E[D].$$

The long run average cost per period for infinite horizon is given by equation (2.2),

$$L(s, S, p) = \frac{K}{1 + R(q)} + cE[D] + hE[X^+] + gE[X^-] - pE[D].$$

Regardless of policy taken, the average per period production amount is always E[D]. According to PA derivatives equations (2.12) and (2.17), we have w.p.1 that

$$(\frac{\partial \overline{x_t^+}}{\partial s})_{IPA} \to P(X > 0) \tag{2.22}$$

$$(\frac{\partial \overline{x_t^-}}{\partial s})_{IPA} \to -P(X < 0) \tag{2.23}$$

$$(\frac{\partial \overline{x_t^+}}{\partial q})_{PA} \to P(X > 0) + P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [E[(s-D)^+] - E[X^+]] \tag{2.24}$$

$$(\frac{\partial \overline{x_t^-}}{\partial q})_{PA} \to -P(X < 0) + P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [E[(s-D)^-] - E[X^-]] \tag{2.25}$$

$$(\frac{\partial \overline{(k * \delta(y_t - x_t)}}{\partial q})_{PA} \to P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [-\frac{k}{1+R(q)}] \tag{2.26}$$

$$(\frac{\partial \overline{((S - x_t) * \delta(y_t - x_t))}}{\partial q})_{PA} \to P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [E[D] - E[D]] = 0 \tag{2.27}$$

$$(\frac{\partial \overline{(-pD_t)}}{\partial q})_{PA} \to P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [-pE[D] + pE[D]] = 0 \tag{2.28}$$

According to Fu's paper, we have

$$\frac{\partial E(X^+)}{\partial s} = P(X > 0) \text{ = eq. (2.22)}$$

$$\frac{\partial E(X^-)}{\partial s} = -P(X < 0) \text{ = eq. (2.23)}$$

$$\frac{\partial E(X^+)}{\partial q} = P(X > 0) + P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [E[(s-D)^+] - E[X^+]] = \text{ eq.(2.24)}$$

$$\frac{\partial E(X^-)}{\partial q} = -P(X < 0) + P(Y = S)E[\frac{f(Z)}{1-F(Z)}] \cdot [E[(s-D)^-] - E[X^-]] \text{ = eq. (2.25)}$$

We also have

$$\frac{\partial(k/(1+R(q)))}{\partial s}=0.$$

$$\frac{\partial(k/(1+R(q)))}{\partial q}=-\frac{r(q)*k}{(1+R(q))^2}=P(Y=S)E[\frac{f(Z)}{1-F(Z)}]\cdot[-\frac{k}{1+R(q)}]=\text{eq. (2.26)}$$

$$\frac{\partial E[D]}{\partial s}=0,$$

$$\frac{\partial E[D]}{\partial q}=0=\text{eq. (2.27).}$$

$$\frac{\partial(-pE[D])}{\partial s}=0,$$

$$\frac{\partial(-pE[D])}{\partial q}=0=\text{eq. (2.28).}$$

Therefore, consistency proof is completed for s and q.

Consistency proof for p is more complicated. First we define N (q) to be the counting process for the demand renewal process:

$$N(q)=\max\{t\mid D_1+D_2+...+D_t\le q\}$$

Then for $t\in M^*(T)$, we have t-t'-1 ~ N (q) and $z_t \sim q-(D_1+D_2\ldots+D_{N(q)})$. Hence,

$$P(t-t'-1=0,z_t=q)=1-F(q),$$

and for n = 1,2,… and $z\in[0,q]$,

$$f_{t-t'-1,z_t}(n,z)=f_n(q-z)*(1-F(z)).$$

Therefore,

$$E[\frac{f(z_t)(t-t')}{1-F(z_t)}] = \frac{(1-F(q))*f(q)}{1-F(q)} + \sum_{n=1}^{\infty}(n+1)*\int_0^q f_n(q-z)*(1-F(z))*\frac{f(z)}{1-F(z)}dz$$

$$= f(q) + \sum_{n=1}^{\infty}(n+1)*\int_0^q f_n(q-z)f(z)dz = f(q) + \sum_{n=1}^{\infty}(n+1)*f_{n+1}(q)$$

$$= \sum_{n=1}^{\infty} nf_n(q)$$

For PA analysis of price, according to equation (2.21), we have w.p.1 that:

i)
$$(\frac{\partial(\overline{k*\delta(y_t - x_t)})}{\partial p})_{PA} = \frac{1}{T+1}\sum_{t\in M^*(T)}[\frac{f(z_t)}{1-F(z_t)}(t-t')a]*[-\frac{k}{1+R(q)}]$$
$$\to \sum_{n=1}^{\infty} nf_n(q)*[-\frac{ak}{(1+R(q))^2}]$$

ii)

$$(\frac{\partial(\overline{(S-x_t)*\delta(y_t - x_t)})}{\partial p})_{PA} = \frac{1}{T}\sum_{t:x_t<s} -(t-t')a + \frac{1}{T+1}\sum_{t\in M^*(T)}[\frac{f(z_t)}{1-F(z_t)}(t-t')a]*[E[D]-E[D]]$$
$$= \frac{1}{T}\sum_{t:x_t<s} -(t-t')a \to \frac{1}{T}*-T*a = -a$$

iii)

$$(\frac{\partial \overline{x_t^+}}{\partial p})_{PA} = \frac{1}{T}[\sum_{t:x_t\geq s, x_{t+1}>0}(t-t')a + \sum_{t:x_{t+1}>0}a] + \frac{1}{T+1}\sum_{t\in M^*(T)}[\frac{f(z_t)}{1-F(z_t)}(t-t')a]*[E[(s-D)^+]-E[X^+]]$$

$$\to \frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)F(s+q-u)du + aP(X>0) + \frac{a\sum_{n=1}^{\infty} nf_n(q)}{1+R(q)}\cdot[E[(s-D)^+ - E[X^+]]$$

The first term on the right hand side is derived as follows:

Take t - t' = n, n=1, 2, …, then we have

$$\frac{1}{T}\sum_{t:x_t\geq s,x_{t+1}>0}(t-t')a \to \frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot P(x_{t'+n}\geq s,x_{t'+n+1}\geq 0)$$

$$=\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot P(D_{t'}+\ldots+D_{t'+n-1}\leq q,D_{t'}+\ldots+D_{t'+n}<s+q)$$

$$=\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot \int_0^q P(D_{t'}+\ldots+D_{t'+n}<s+q \mid D_{t'}+\ldots+D_{t'+n-1}=u)f_n(u)du$$

$$=\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot \int_0^q P(D_{t'+n}<s+q-u)f_n(u)du$$

$$=\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)F(s+q-u)du$$

iv)
$$(\frac{\partial \overline{x_t^-}}{\partial p})_{PA}=\frac{1}{T}[\sum_{t:x_t\geq s,x_{t+1}<0}-(t-t')a-\sum_{t:x_{t+1<0}}a]\frac{1}{T+1}\sum_{t\in M^*(T)}[\frac{f(z_t)}{1-F(z_t)}(t-t')a]*[E[(s-D)^-]-E[X^-]]$$

$$\to -\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)(1-F(s+q-u))du-aP(X<0)+\frac{a\sum_{n=1}^{\infty} nf_n(q)}{1+R(q)}[E[(s-D)^-]-E[X^-]]$$

The first term on the right hand side is given by the similar derivation as iii)

$$\frac{1}{T}\sum_{t:x_t\geq s,x_{t+1}<0}-(t-t')a \to \frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot P(x_{t'+n}\geq s,x_{t'+n+1}<0)$$

$$=-\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot P(D_{t'}+\ldots+D_{t'+n-1}\leq q,D_{t'}+\ldots+D_{t'+n}>s+q)$$

$$=-\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot \int_0^q P(D_{t'}+\ldots+D_{t'+n}>s+q \mid D_{t'}+\ldots+D_{t'+n-1}=u)f_n(u)du$$

$$=-\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\cdot \int_0^q P(D_{t'+n}>s+q-u)f_n(u)du$$

$$=-\frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)(1-F(s+q-u))du$$

v)
$$(\frac{\partial(\overline{-pD_t})}{\partial p})_{PA}=-\frac{1}{T}\sum_{t=1}^{T}(-2ap+b+\varepsilon_t)+\frac{1}{T+1}\sum_{t\in M^*(T)}[\frac{f(z_t)}{1-F(z_t)}(t-t')a]*[-pE[D]+pE[D]]$$

$$\to 2ap-b$$

For (i), we have

$$F_n(q) = P(D_1 + \ldots + D_n \le q)$$
$$= P(\varepsilon_1 + \ldots + \varepsilon_n \le q - n(b - ap))$$
$$= F_{\varepsilon,n}(q - n(b - ap))$$

,

and

$$\frac{\partial F_n(q)}{\partial p} = \frac{\partial F_{\varepsilon,n}(q - n(b - ap))}{\partial p}$$
$$= f_{\varepsilon,n}(q - n(b - ap)) * na = a * nf_n(q)$$

Hence,

$$\frac{\partial (k/(1 + R(q)))}{\partial p} = -\frac{k}{(1 + R(q))^2}\frac{\partial R(q)}{\partial p}$$
$$= -\frac{k}{(1 + R(q))^2}\frac{\partial \sum_{n=1}^{\infty} F_n(q)}{\partial p}$$
$$= -\frac{ak\sum_{n=1}^{\infty} nf_n(q)}{(1 + R(q))^2}$$

This is consisted with limitation developed in part i).

For (ii), we have $\frac{\partial (E[D])}{\partial p} = -a$, which is consisted with result from ii).

For (iii), applying the similar approach by Fu (1994a), we have

$$E[X^+] = \frac{1}{1 + R(q)}[\overline{F}(s + q) + \overline{F}(s)R(q) + \int_0^q R(u)F(s + q - u)du],$$

and $P(X > 0) = \frac{1}{1 + R(q)}[F(s + q) + F(s)R(q) + \int_0^q R(u)f(s + q - u)du$

Where $\overline{F}(x) = \int_0^x F(x)dx$.

Differentiate the expected value with respect to price, we have

$$\frac{\partial E[X^+]}{\partial p} = \frac{1}{1+R(q)}[aF(s+q) + aF(s)R(q) + \overline{F}(s)a\sum_{n=1}^{\infty} nf_n(q) + a\int_0^q R(u)f(s+q-u)du$$
$$+ a\int_0^q \sum_{n=1}^{\infty} nf_n(u)\cdot F(s+q-u)du - a\sum_{n=1}^{\infty} nf_n(q) * E[X^+]]$$
$$= \frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)F(s+q-u)du + aP(X>0) + \frac{a\sum_{n=1}^{\infty} nf_n(q)}{1+R(q)}\cdot[E[(s-D)^+ - E[X^+]$$

This is also consisted with limitation derived in part iii).

For (iv)

$$\frac{\partial E[X^-]}{\partial p} = \frac{\partial E[X^+]}{\partial p} - \frac{\partial E[X]}{\partial p}$$
$$= \frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)F(s+q-u)du + aP(X>0) + \frac{a\sum_{n=1}^{\infty} nf_n(q)}{1+R(q)}\cdot[E[(s-D)^+] - E[X^+]]$$
$$+ \frac{a\sum_{n=1}^{\infty} nf_n(q)}{(1+R(q))^2}\cdot(q+\overline{R}(q)) - \frac{a\sum_{n=1}^{\infty} nF_n(q)}{1+R(q)} - a$$
$$= -aP(X<0) + \frac{a}{1+R(q)}\{\sum_{n=1}^{\infty} n\int_0^q f_n(u)F(s+q-u)du - \sum_{n=1}^{\infty} nF_n(q)\}$$
$$+ \frac{a\sum_{n=1}^{\infty} nf_n(q)}{1+R(q)}\{E[(s-D)^+] - E[X^+] + E[X] - s + E[D]\}$$
$$= -aP(X<0) - \frac{a}{1+R(q)}\sum_{n=1}^{\infty} n\int_0^q f_n(u)(1-F(s+q-u))du$$
$$+ \frac{a\sum_{n=1}^{\infty} nf_n(q)}{1+R(q)}\{E[(s-D)^-] - E[X^-]\}$$

So it's consisted with limitation in part iv).

For v), we have

$\frac{\partial(-pE[D])}{\partial p} = \frac{\partial(-p(b-ap))}{\partial p} = 2ap - b$, which is equal to the result from part v).

Therefore, consistency proof is completed for PA estimator of price p.

2.5 Optimization Example

In this section, the sample path derivatives derived from previous sections are used in the RMSA algorithm to find the optimal setting of the (s, S, p) system. The algorithm performance obtained from RMSA is compared with those from SPSA and SAN approaches in some numerical cases.

2.5.1 Robbins-Monro Stochastic Approximation (RMSA)

Our goal is to find θ* that solves $\min_{\theta \in C} L(\theta)$, where C represents a constraint set defining the allowable values for the parameters θ. We are interested in the gradient vector of the loss function with respect to the parameters:

$$L'(\theta) = \begin{bmatrix} \frac{\partial L}{\partial s} \\ \frac{\partial L}{\partial q} \\ \frac{\partial L}{\partial p} \end{bmatrix}.$$

Then for local optimization, a necessary condition for optimization when L is continuously differentiable is that θ* satisfies: L'(θ*) = 0. Using a Robbins-Monro stochastic approximation algorithm (Kushner and Yin 1997) and the gradient estimator derived in the previous section, we apply the following iterative gradient-based procedure to update the parameter values, in order to reach a local optimum:

$$\begin{bmatrix} s_{k+1} \\ q_{k+1} \\ p_{k+1} \end{bmatrix} = \begin{bmatrix} s_k \\ q_k \\ p_k \end{bmatrix} - e_k \begin{bmatrix} (\frac{\partial L}{\partial s})_{PA} \\ (\frac{\partial L}{\partial q})_{PA} \\ (\frac{\partial L}{\partial p})_{PA} \end{bmatrix}_{s_k, q_k, p_k}, \tag{2.29}$$

where k is the iteration number, and the gain sequence $\{e_k\}$ is a version of the accelerated harmonic series given by e/E_k, where e is some constant and

$$E_{k+1} = \begin{cases} E_k + 1 & \text{if } \operatorname{sgn}(\partial L/\partial s)_{PA,k+1} \neq \operatorname{sgn}(\partial L/\partial s)_{PA,k} \\ & \quad \& \operatorname{sgn}(\partial L/\partial q)_{PA,k+1} \neq \operatorname{sgn}(\partial L/\partial q)_{PA,k} \\ & \quad \& \operatorname{sgn}(\partial L/\partial p)_{PA,k+1} \neq \operatorname{sgn}(\partial L/\partial p)_{PA,k} \\ E_k & \text{otherwise} \end{cases}$$

with $E_0 = 1$, and *sgn* is the sign function of a vector of parameters. The step size changes only when all three signs of the vector elements change simultaneously. Since s, q, and p must be positive, we project back to the previous point whenever the algorithm brings s, q or p negative. The values of s, q and p are updated every T periods, with the PA estimator reinitialized at each update. Furthermore, we take the starting point to be $s_0 = q_0 = E[D]/2$, and $p_0 = (p_{min} + p_{max})/2$, where p_{min} and p_{max} are the lower bound and upper bound of price range, respectively. We expect that the parameters e and T greatly affect the initial convergence rate of the algorithm.

2.5.2 Simultaneous Perturbation Stochastic Approximation (SPSA)

SPSA is one of the gradient-free methods, which does not depend on direct gradient information or measurements. It is based on an approximation to the gradient formed from measurements of the loss function. That is to say, such algorithm do not require

the detailed knowledge of the functional relationship between the parameters being optimized and the loss function being minimized, which is required in gradient-based algorithms. Spall (1998) illustrated SPSA for efficient optimization.

The recursive procedure we consider in the SPSA is the same as RMSA, i.e. equation (2.29). All elements of θ_k are randomly perturbed together to obtain two measurements L(.), but each component of estimated $L'_k(\theta_k)$ is formed from a ratio involving the individual components in the perturbation vector and the difference in the two corresponding measurements. Thus we have

$$L'_k(\theta_k) = \frac{L(\theta_k + r_k\Delta_k) - L(\theta_k - r_k\Delta_k)}{2r_k} \begin{bmatrix} \Delta_{k1}^{-1} \\ \Delta_{k2}^{-1} \\ \Delta_{k3}^{-1} \end{bmatrix} \tag{2.30}$$

We pick $e_k = e/(E+k)^{\alpha}$, $r_k = r/k^{\beta}$, where e and E is some constants controlling algorithm's convergence speed. A simple choice for each component of Δ_k is to use a Bernoulli ±1 distribution with probability of ½ for each ±1 outcome. We take the starting point at $s_0 = q_0 = E[D]/2$, and $p_0 = (p_{min} + p_{max})/2$.

2.5.3 Simulated Annealing (SAN)

RMSA and SPSA only guarantee to yield a local optimum, while SAN aims to find a global solution from among multiple local solutions. The general sequence of steps is as following:

Algorithm SAN:

Step1: Choose an initial A_0 and set of current parameter values θ_{curr}; determine $L(\theta_{curr})$.

Step2: Randomly determine a new value of θ, θ_{new}, that is "close" to the current value, and determine L (θ_{new}).

Step 3: Compare the two L values: Let $\delta = L(\theta_{new}) - L(\theta_{curr})$. Accept θ_{new} if $\delta < 0$. Alternatively, if $\delta >= 0$, accept the new point θ_{new} only if a uniform (0, 1) random variable U satisfies $U <= \exp(-\delta/A)$.

Step 4: Repeat steps 2 and 3 for some period until either the budget of function evaluations allocated for that A has been used or the system reaches some state of equilibrium.

Step 5: Lower A according to the annealing schedule $A_k = A_0/\ln(1+k)$, and return to step 2. Continue the process until the total budget for function evaluations has been used or some indication of convergence is satisfied.

2.5.4 Numerical Analysis

Our numerical study is based on the data collected from a specialty retailer of high-end women's apparel (Federgruen and Heching 1999). Table 2.1 summarizes all parameters for the base scenarios pertaining to the dress. The variables ε_t in the

additive stochastic demand function (equation (2.1)) are normally distributed with zero mean and standard deviation $\delta(p)*cv$, where cv a specified coefficient of variation. The demand thus follows a normal distribution with mean of $\delta(p) = (b-ap)$ and standard deviation of $cv*(b-ap)$. We truncated at $-\delta(p)$ to preclude negative demand realizations.

Table 2.1 Base Parameters for Dress.

Item	b	a	cv	k	c	h	g	Price Range
Dress	174	3	0.25	0	22	0.22	21.78	25-44

In our numerical experiment, initial values of the parameters are set at (s, q, p) = (35, 35, 34.5). We choose the update period T = 100 and run 100 replications. Each simulation replication for RMSA and SPSA is terminated when the sum of the gradient estimate components for the three parameters is less than 0.001 or the number of iteration is greater than 10,000, while SAN simulation terminates when the difference of loss value between two consecutive iterations is less than 0.1. Choosing proper gain sequence coefficients might be the most difficult issue in SPSA, we fix $\alpha=1$, $r=0.5$, $\beta=0.2$, and we choose step size e and E by trial-and-error.

We consider three cases corresponding to different fixed cost k and holding cost h with cv = 0.25:

Case 1 – k = 0, h = 0.22;

Case 2 – k = 0, h = 5;

Case 3 – k = 100, h = 0.22.

Tables 2.2, 2.3, and 2.4 show 95% confidence intervals of optimal values for three cases by RMSA, SPSA, and SAN, respectively. We found that optimal price and loss is close for three algorithms and RMSA has relatively large variance. When fixed ordering cost is 0 in case 1, the higher s we have, the least possibility would shortage occur. For case 2, the results indicate that s and S decreases as holding cost increases. Reorder point S converges closely for three algorithms, which is significantly different from the case of h=0.22. As expected, loss function becomes more sensitive to S when holding cost increases. The lower S we have, the less would holding cost be. Notice that fixed ordering cost is zero in this case, therefore, number of ordering period is not a main concern. We also find that the variance is generally small in this case since stock level (s, S) affects loss function more here. There is tradeoff between ordering cost and holding cost when fixed ordering cost is not zero. As we already known from previous case, high holding cost leads to low S. However, low S leads to frequent ordering, consequently, causing large ordering cost if ordering setup cost exists. In case 3, where holding cost is small compared with fixed ordering cost, q is relatively large, decreasing the number of ordering cycles.

The optimal value of q is small for zero fixed ordering cost, substantially smaller than expected demand. In this situation, holding costs dominate, since frequent ordering is not penalized. Comparing case 1 and case 2, we find that s and S decrease as holding cost increases. In all three cases, price doesn't differ much, i.e., price does not appear to be a major determinant for the various inventory-related cost scenarios, which is also consistent with our PA gradient estimators.

Table 2.2 95% Confidence Interval for cv = 0.25 by RMSA.

95% C. I.	K = 0, h = 0.22	K = 0, h = 5	K = 100, h = 0.22
s	65.92 ± 1.74	50.13 ± 0.13	68.55 ± 4.03
S	95.05 ± 6.14	67.80 ± 0.65	181.57 ± 9.72
P	39.17 ± 0.39	39.84 ± 0.07	39.29 ± 0.74
L	-980.14 ± 9.57	-895.42 ± 6.55	-926.01 ± 7.36

Table 2.3 95% Confidence Interval for cv = 0.25 by SPSA.

95% C. I.	K = 0, h = 0.22	K = 0, h = 5	K = 100, h = 0.22
s	42.13±0.33	35.19±0.31	68.29±1.61
S	75.85±1.20	66.67±0.46	138.27±1.55
P	40.37±0.28	40.67±0.08	40.15±0.08
L	-973.33±6.75	-890.23±6.50	-936.20±7.37

Table 2.4 95% Confidence Interval for cv = 0.25 by SAN.

95% C. I.	K = 0, h = 0.22	K = 0, h = 5	K = 100, h = 0.22
s	60.94±1.34	38.98±0.73	1.47
S	88.86±3.61	64.46±0.53	139.87±13.92
P	39.99±0.02	40.86±0.04	39.87±0.05
L	-993.74±5.59	-898.88±6.17	-935.68±11.22

Figures 2.4, 2.5, 2.6, and 2.7 illustrate the convergence rate of RMSA algorithm for three cases based on one common run for L, p, s and S, respectively. The figures show that the algorithm converges very fast at the beginning. The fluctuations in case 3 are due to a large initial step size.

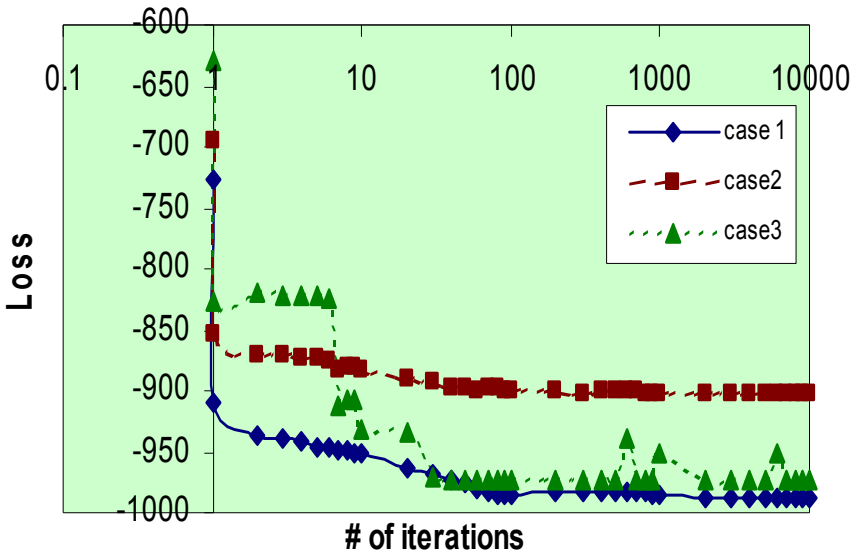

Fig. 2.4 Expected loss as function of iteration number by RMSA

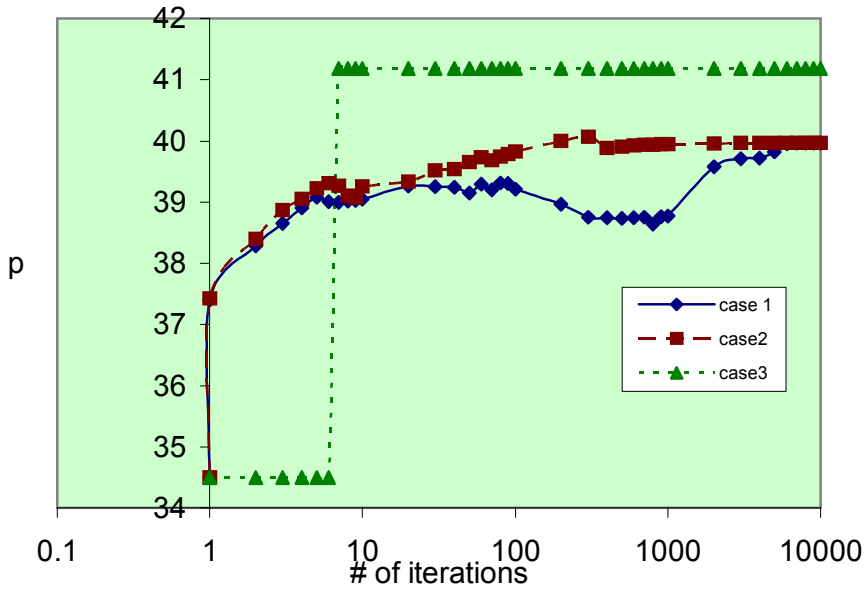

Fig. 2.5 Selling price as function of iteration number by RMSA

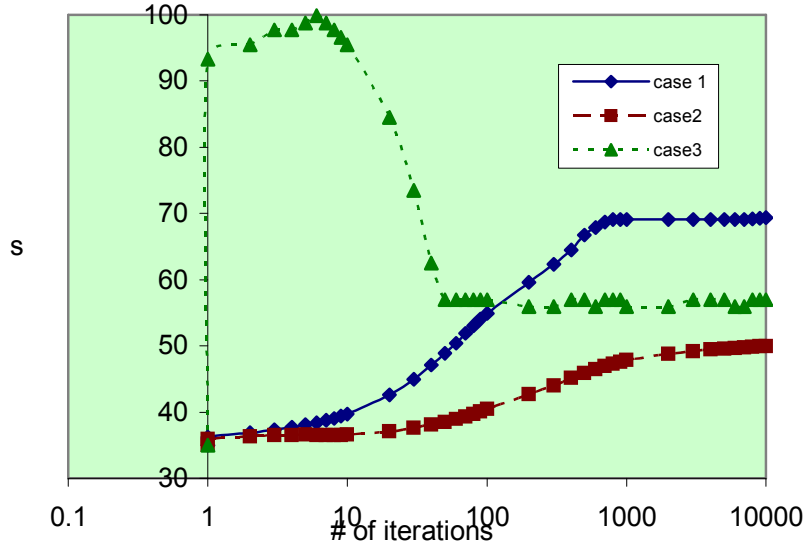

Fig.2.6 Base stock level as function of iteration number by RMSA

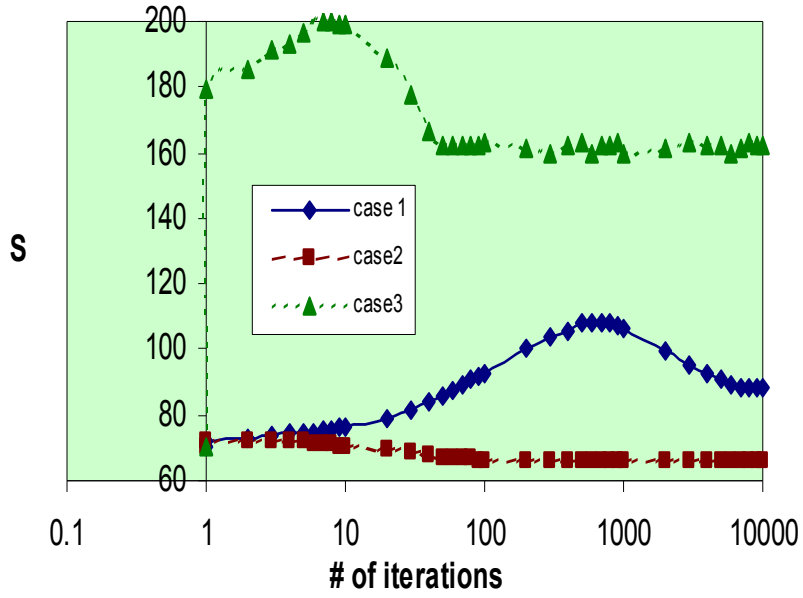

Fig.2.7 Reorder stock level as function of iteration number by RMSA

Figures 2.8, 2.9, 2.10 and 2.11 compare convergence rate of three algorithms in case 1 based on one common random number run for L, p, s and S, respectively. From Figures 2.8 and 2.9 we found RMSA has the greatest convergence rate, at least converges very fast at the beginning. Loss and price have a considerable jump at the first iteration. RMSA is efficient because it relies on direct gradient estimator. The figures illustrate strong dependence between loss and price: loss converges to the same level in spite of the fact that S and s don't approach to close values from three algorithms. The basic idea behind SAN is to randomly pick the next step, and it allows worse objective value exist, so it's not surprise when we observe that SAN converges slower than RMSA and SPSA for the first 100 iterations.

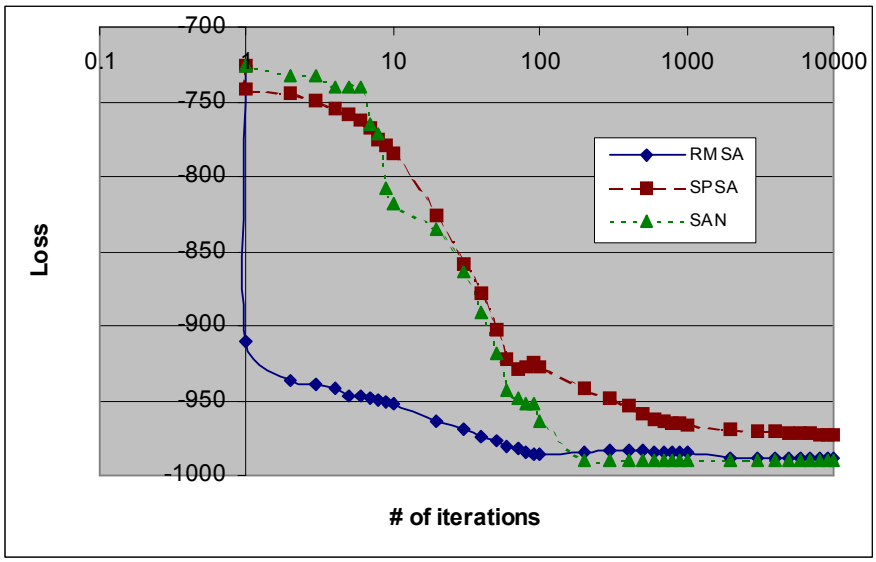

Figure 2.8 Expected loss as function of iteration number – case 1

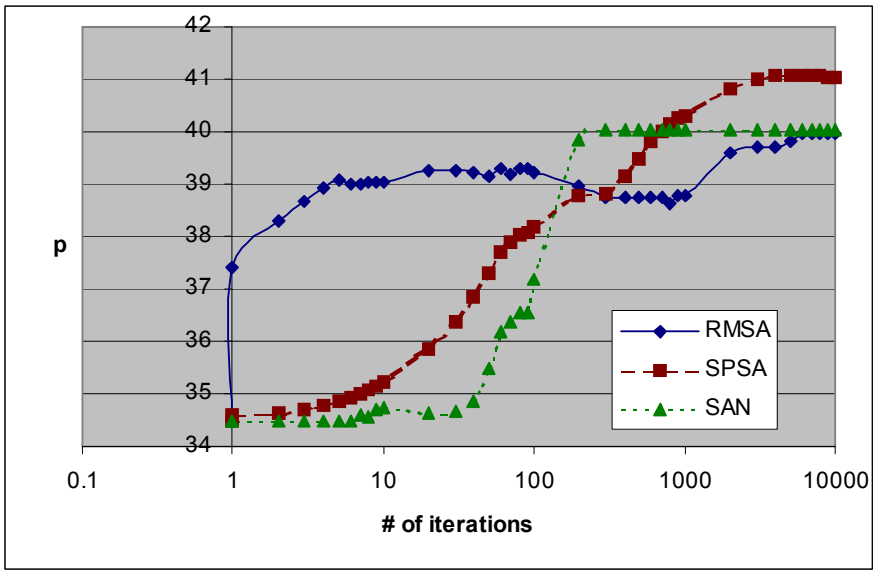

Figure 2.9 Selling price as function of iteration number – case 1

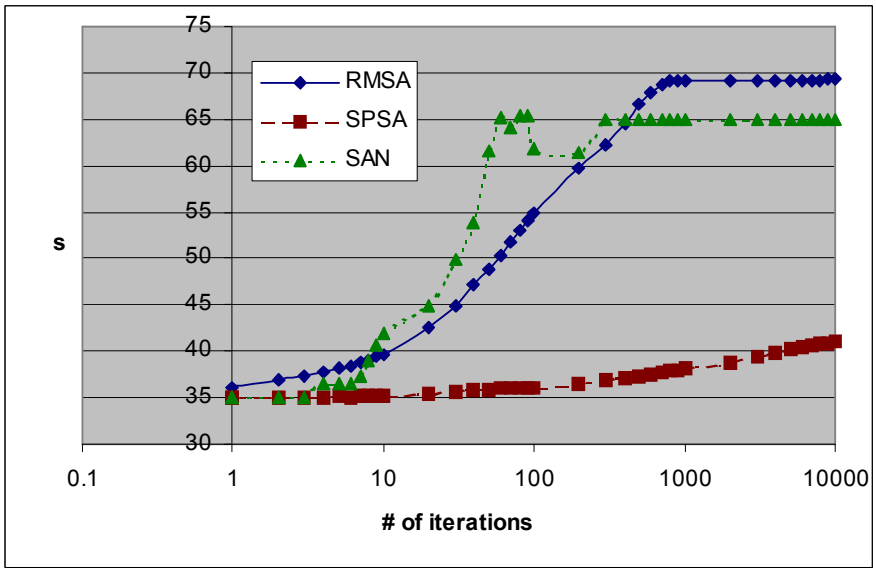

Figure 2.10 Base stock level as function of iteration number – case 1

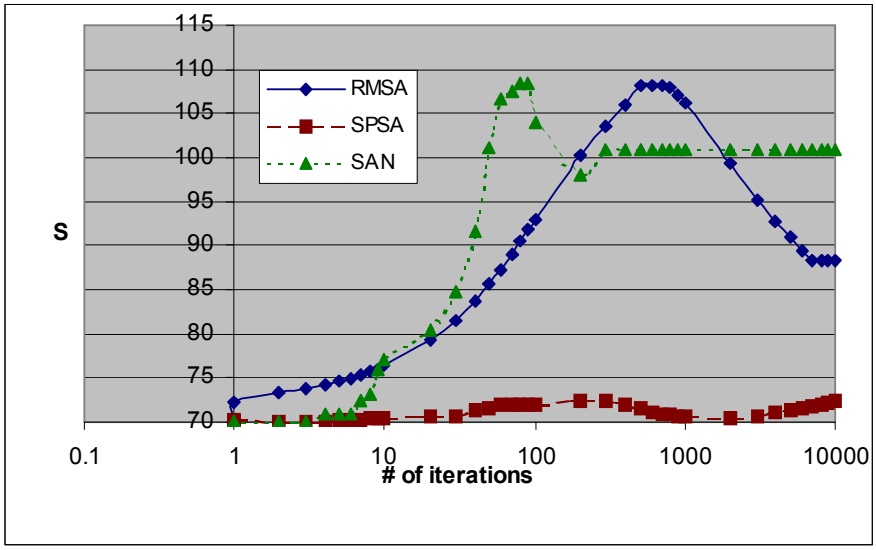

Figure 2.11 Reorder stock level as function of iteration number – case 1

Figures 2.12-2.15 illustrate relationship between parameters, objective function and number of iterations in case 2. RMSA is the most efficient algorithm among three stochastic optimization algorithms due to PA derivatives. Figures 2.16-2.19 show the convergence rate for L, p, s and S in case 3. The fluctuation on RMSA and SPSA is due to the large step size we choose to jump out initial trough.

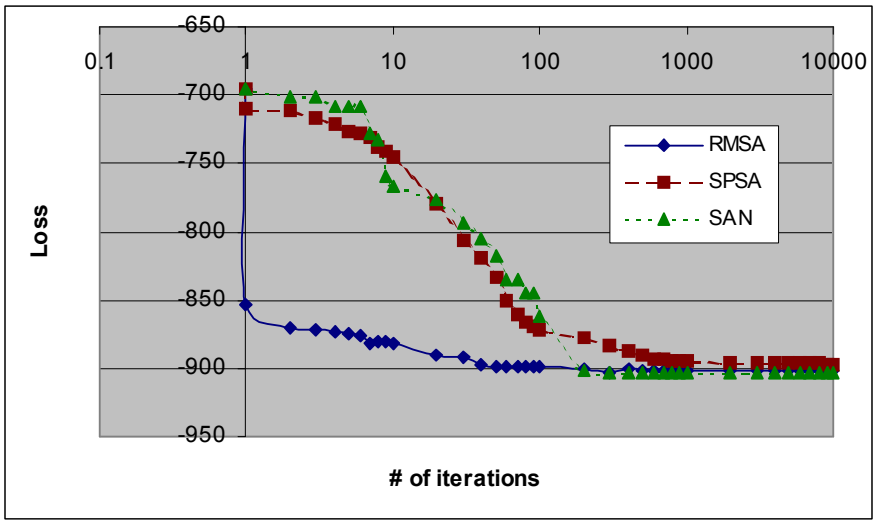

Figure 2.12 Expected loss as function of iteration number – case 2

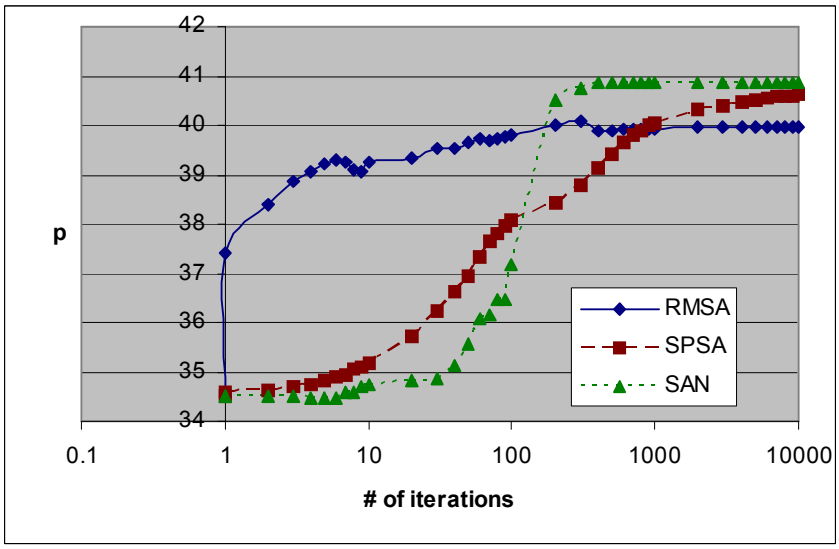

Figure 2.13 Selling price as function of iteration number – case 2

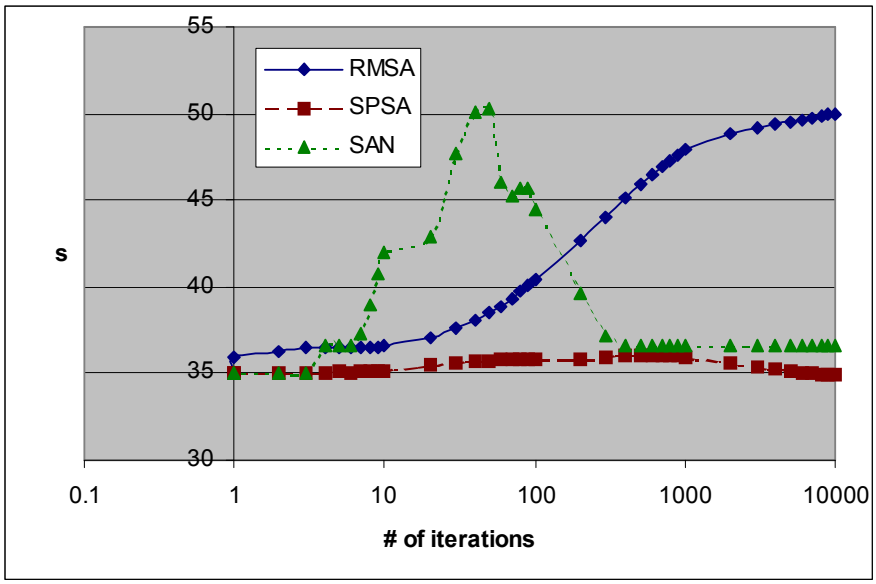

Figure 2.14 Base stock level as function of iteration number – case 2

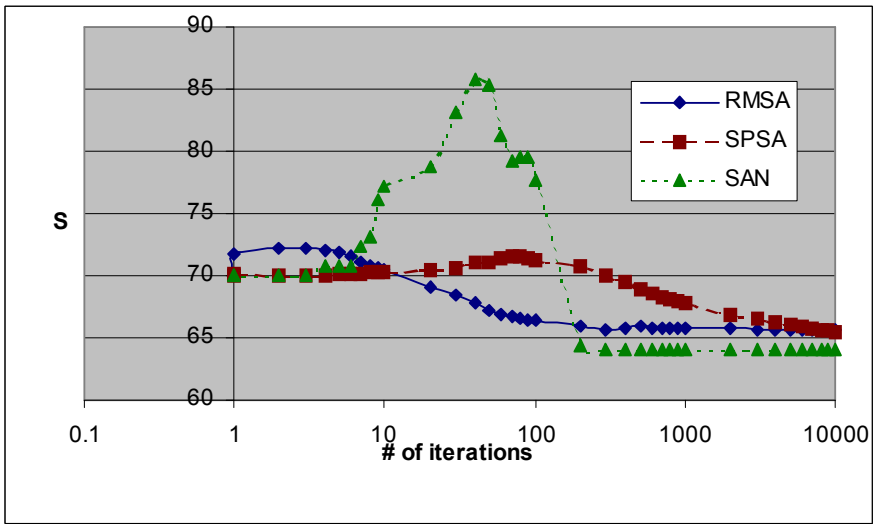

Figure 2.15 Reorder stock level as function of iteration number – case2

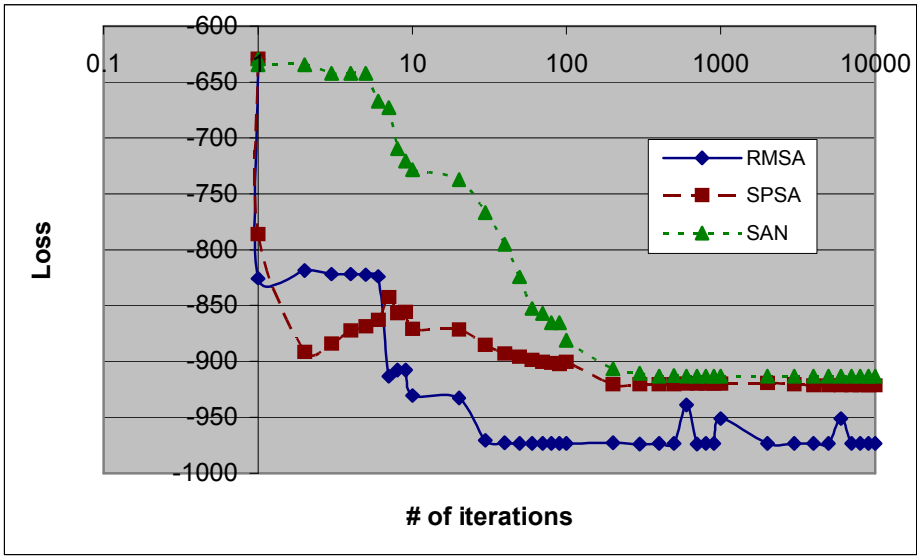

Figure 2.16 Expected loss as function of iteration number – case 3

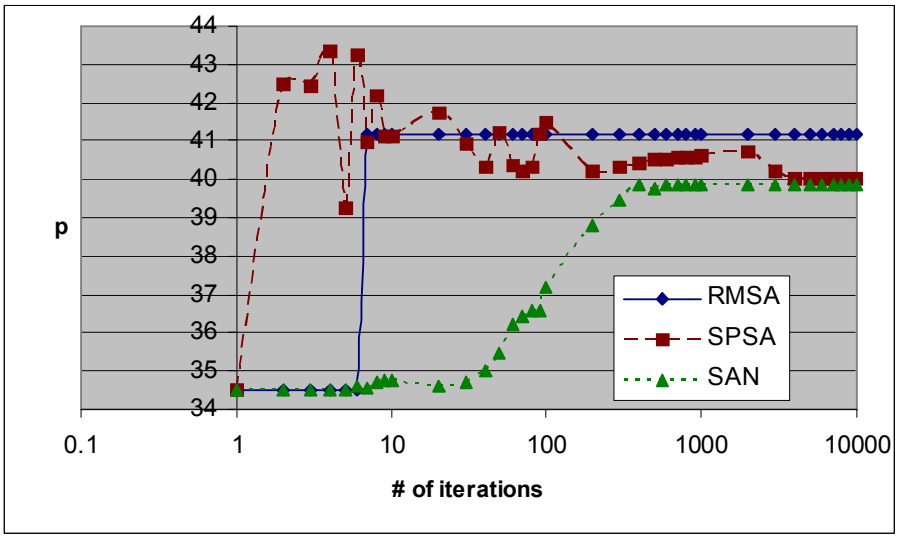

Figure 2.17 Selling price as function of iteration number – case 3

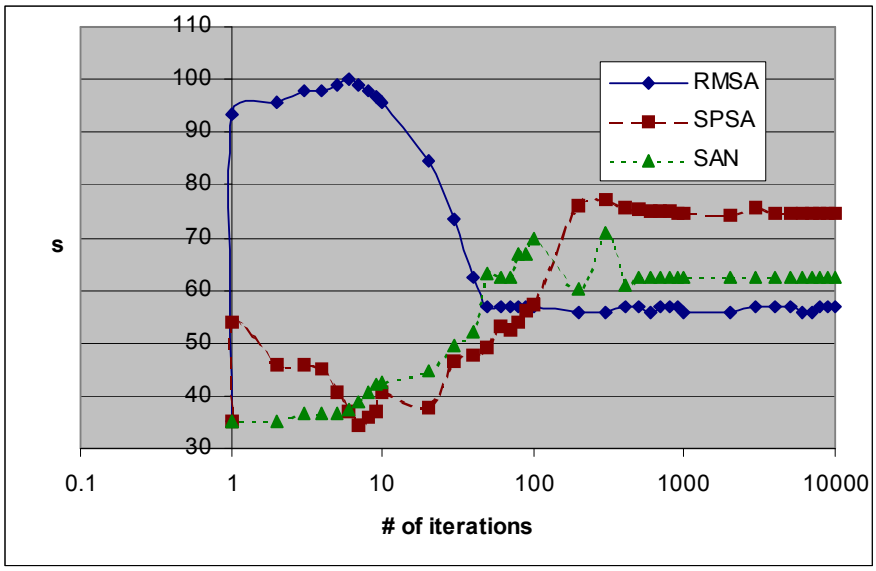

Figure 2.18 Base stock level as function of iteration number – case 3

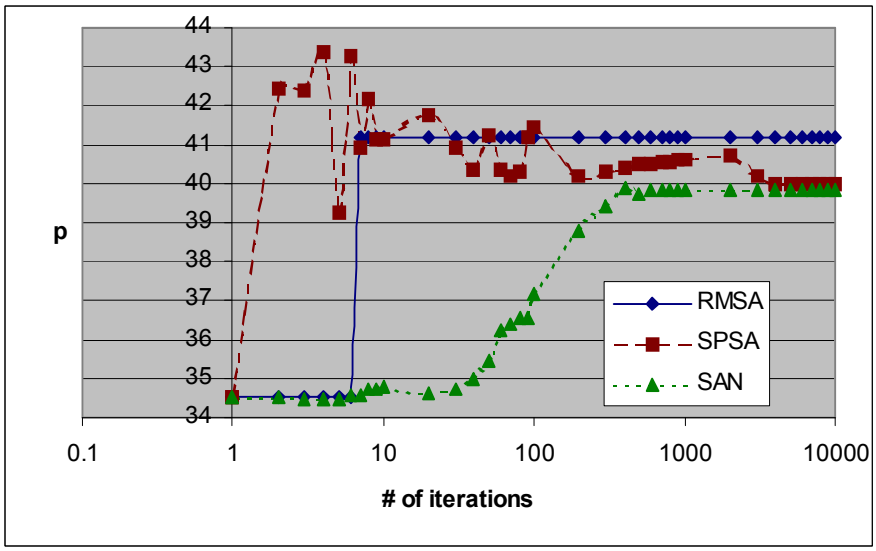

Figure 2.19 Reorder stock level as function of iteration number – case 3

Finally, we investigate the impact of price elasticity of demand by modifying the slope a of the demand function to a = 1, a = 3, and a = 5, with k = 100 and h = 0.22. Table 2.5 shows the values of the control parameters and total loss. Price decreases dramatically as a increases. Price elasticity measures the change of demand to change of price; thus, when it goes up, a manufacturer has to reduce price to attract more consumers so as to increase revenue.

Table 2.5 Experimental Results for Different Price Elasticity of Demand a

(cv=0.25, k=100, h=0.22)

		RMSA		
a	s	S	p	L
1	93.53	166.11	43.91	-2811.44
3	53.70	133.32	40.37	-961.37
5	47.62	121.07	25.04	-123.23
		SPSA		
a	s	S	p	L
1	353.10	717.89	43.44	-2824.00
3	74.68	150.84	39.99	-921.82
5	34.59	82.81	28.17	-173.57
		SAN		
a	s	S	p	L
1	157.08	244.28	44.00	-2821.64
3	62.40	118.35	39.84	-913.14
5	41.21	168.27	28.27	-203.92

Some general observations from the simulation results:

1. T = 100 is not sufficient for reaching steady state for large fixed ordering cost, since iterate updates are not carried out at regenerative points, so we have the `last period effect'.

2. In some simulations, periodic behavior occurs in the iterations, due to the implementation of the gain sequence, which only decreases if all three components in the gradient change signs.

3. Adjusting step size coefficients is critical and depends on initial conditions. We use trial-and-error in the simulation. All stochastic optimization techniques suffer from this problem.

2.6 Conclusions

This essay presents a period review inventory model with price-dependent uncertain demand. The proposed inventory control policy reflects a common practice in some industries. To minimize the expected loss, management determines both the optimal stock level and price. Using perturbation analysis, we develop sample path derivatives for this (s, S, p) inventory model, which could be incorporated into gradient-based algorithms to select optimal values for the three controllable parameters. Consistency proofs are provided for the infinite horizon case. Some numerical results for simulation optimization are presented using a Robbins-Monro stochastic approximation algorithm. We also apply two other stochastic optimization algorithms to compare with RMSA in different scenarios. Though we don't have analytical outcomes to replicate in the stochastic demand circumstance, the results from three algorithms are comparable and explainable. Measuring of convergence rate for three algorithms demonstrate that RMSA is most efficient, at least at the beginning of the simulation.

We study the effect of unit holding cost, fixed ordering cost and price elasticity of demand for the (s, S, p) system. The results are intuitive and verify our algorithms to some extent. When holding cost is large compared to fixed ordering cost, reorder stock level S will decrease to reduce holding cost. On the contrary, S and q will increase so that number of ordering period decreases. Price goes down as price elasticity of demand increases. We also found that selling price is the determinant factor.

For future research, useful extensions of our model include applying our method to stochastic lead-time scenarios, and investigating a multiple-market setting problem in which demand distribution is dependent on market selection.

Chapter 3

Applying Model Reference Adaptive Search to American-style Option Pricing

This essay is organized as follows. The problem setting is described in Section 1. The MRAS algorithm applied to American-style option pricing is discussed in Section 2, and it is implemented in pricing American-style call options with dividend, pricing American-style put options written on underlying assets following geometric Brownian motion and Merton jump diffusion model, and pricing American-style Asian call options without dividend in Section 3. The results from MRAS algorithm are compared with CE method in this section. Finally we offer some conclusions based on the numerical results in Section 4.

3.1 Problem Setting

We consider the American option pricing problem as a maximization problem and apply optimization techniques to parameterize the early exercise boundary. The value of American call option written on a single stock with finite early exercise dates can be written as

$$\max_{S_i^*} E(L(S_i^*)],$$

where

$$L(S_i^*) = \sum_{i=0}^{n-1} 1\{\bigcap_{j=0}^{i-1} S_j < S_j^*, S_i > S_i^*\}(S_i - K)e^{-rt_i} + 1\{S_1 < S_1^*, \ldots S_{n-1} < S_{n-1}^*\}(S_T - K)^+ e^{-rT}, \quad (3.1)$$

$1\{\bullet\}$: indicator function,

K: strike price,

r: risk free rate,

T: maturity,

n: number of exercise opportunities, including the exercise at maturity,

S_i^*: early exercise threshold at exercise date t_i; the parameters to be estimated for the optimization problem. Note that for convenience we will use S* to represent the critical prices set $\{S_i^*\}$ in the following work,

S_i: stock price at exercise date t_i,

L: the sample performance is the net present value of the option payoff.

The first term on the right side is the payoff of early exercise, and the second term is the payoff without early exercise, i.e., the payoff at the time of maturity.

Similarly, the American put option can be written as

$$\max_{S_i^*} E(L(S_i^*)]$$

and

$$L(S_i^*) = \sum_{i=0}^{n-1} 1\{\bigcap_{j=0}^{i-1} S_j > S_j^*, S_i < S_i^*\}(K - S_i)e^{-rt_i} + 1\{S_1 > S_1^*, \ldots S_{n-1} > S_{n-1}^*\}(K - S_T)^+ e^{-rT}. \quad (3.2)$$

For convenience, we omit subscript i of S_i^*, using S* to represent the set of critical prices in the following. Once we find the estimates for the thresholds at all exercise points through optimization, we can obtain the value of the option through a forward moving simulation starting from time 0. The procedure simultaneously optimizes all parameters by iteration, and no dynamic programming is involved. In addition, this flexible value function can handle pure-jump and jump-diffusion processes, which can sometimes be problematic for the most popular pricing methods, such as partial differential equation methods, binomial trees, and other lattice methods. In the following numerical examples, we consider the underlying asset following two stochastic processes – geometric Brownian motion and jump diffusion model from Merton (1976).

3.2 Algorithm Description

MRAS is an adaptive algorithm equipped with a random mechanism and a reference model, which work with a family of parameterized distributions on the solution space. The basic idea is to assign more weight to the solutions that have better performance at each step. Kullback-Leibler (KL) divergence is a natural distance measure between two probability distributions in probability theory. At each iteration, samples are generated according to the distribution that has the minimum KL-divergence with respect to the reference model from the previous iteration, and the parameters of the

next distribution are updated based on those samples in a way so that the distribution possesses the minimum KL-divergence with respect to the current reference model.

The main difference between MRAS and CE is that CE method uses a single optimal (importance sampling) distribution focused on the set of optimal solutions (i.e., zero variance) to guide the updating of parameters, while the MRAS uses a sequence of intermediate reference distributions to facilitate and direct its parameter updating associated with the family of parameterized distributions during the search process. We will compare the results from MRAS with those from CE methods in the following sections.

The MRAS method also resembles another model-based method - the estimation of distribution algorithms (EDAs). EDAs were introduced in the field of Evolutionary Computation for the first time by Mühlenbein and Paaβ (1996). EDAs generate new solutions according to the probability distribution of all promising solutions of the previous iteration. No recombination process such as crossover and mutation operators is involved to avoid the disruptions of partial solutions of genetic algorithm. In EDAs the problem specific interactions among the variables of individuals are taken into consideration and the interrelations are expressed explicitly through the joint probability distribution associated with the individuals of variables selected at each generation. New population is generated by sampling the probability distribution, which is estimated from a database containing selected individuals of the previous generation. Larranaga et al. (1999) and Paul et al. (2002) give reviews of implementing EDA approaches using various underlying probabilistic models. However, the estimation of the joint probability distribution associated with the

selected samples is a bottleneck of this method. There is no easy method to calculate it. If the distribution is more general, we expect to get high quality result, but calculation of this distribution is time consuming and complicated, and sampling of new instances using this distribution is not an easy task. In contrast, MRAS uses the sequence of reference models implicitly to guide the parameter updating procedure and there is no need to calculate them explicitly; therefore MRAS overcomes the most difficult obstacle of the EDAs.

Hu et al. (2005) demonstrate the global convergence of MRAS for a class of parameterized probability distributions called Natural Exponential Family (NEF), which includes multivariate normal distribution. In the following numerical experiments, we assume the parameters to be estimated are multi-normally distributed with p.d.f. of

$$f(S^*, \mu_k, \Sigma_{k+1}) = \frac{1}{\sqrt{(2\pi)^n \mid \Sigma_k \mid}} \exp(-\frac{1}{2}(S^* - \mu_k)^T \Sigma_k^{-1} (S^* - \mu_k)) \tag{3.3}$$

where μ_k is the mean vector and Σ_k the covariance matrix at iteration k, and the parameters that minimize the KL divergence are updated as

$$\mu_{k+1} = \frac{\sum \{[U(L(S^*))^k / f(S^*, \mu_{k+1}, \Sigma_{k+1})] I_{\{L(S^*) \geq \bar{\gamma}_{k+1}\}} S^*\}}{\sum \{[U(L(S^*))^k / f(S^*, \mu_{k+1}, \Sigma_{k+1})] I_{\{L(S^*) \geq \bar{\gamma}_{k+1}\}}\}}, \tag{3.4}$$

and

$$\Sigma_{k+1} = \frac{\sum \{[U(L(S^*))^k / f(S^*, \mu_{k+1}, \Sigma_{k+1})] I_{\{L(S^*) \geq \bar{\gamma}_{k+1}\}} (S^* - \mu_{k+1})(S^* - \mu_{k+1})^T\}}{\sum \{[U(L(S^*))^k / f(S^*, \mu_{k+1}, \Sigma_{k+1})] I_{\{L(S^*) \geq \bar{\gamma}_{k+1}\}}\}}, \tag{3.5}$$

where the function U (·) is used to account for the case that the values of L(S*) are negative.

We employed a method called acceptance-rejection to the optimization problem. The key idea is to generate a random vector from a parameterized distribution (multivariate normal distribution in our cases). The sample is accepted or rejected depending on whether it falls or not in the interval of interest. The accepted sample can be viewed as the one generated from the truncated multivariate normal distribution. In the case of pricing American-style put options, the critical price increases as time approaches maturity, and the critical price at the maturity is the strike price K. We generate the critical price increments at the exercise dates from a multivariate normal distribution with given parameters. For those increments not at the first exercisable date, we accept the positive ones and rule out the negative ones. In addition, we only accept those samples in which the critical price at the last exercise date before maturity is less than the strike price K, i.e. the critical price at maturity. Similarly, in the case of American-style call options: we accept the random samples that give negative increments at the exercisable dates except the first date, and satisfy the constraint that the threshold at the last exercise date before maturity is larger than the strike price K.

To avoid the (possible) premature convergence to a degenerate distribution and result in a sub-optimal solution, we applied a dynamic smoothing scheme as described in Kroese et al. (2004) instead of a fixed scheme. Define

$$\beta_k = \beta - \beta(1-\frac{1}{k})^q, \tag{3.6}$$

and the smoothed parameter updating procedure is

$$\hat{\mu}_k = \beta_k \mu_k + (1-\beta_k)\hat{\mu}_{k-1}, \tag{3.7}$$

and

$$\hat{\Sigma}_k = \beta_k \Sigma_k + (1-\beta_k)\hat{\Sigma}_{k-1}. \tag{3.8}$$

where k is the iteration number, β is a smoothing constant (0.8 in our examples), and q is an integer (5 in our examples).

In practice, the MRAS algorithm in the Monte Carlo version works as follows:

Algorithm MRAS

---

1. Initialize: quantile parameter ρ_0, initial sample size N_0, the multivariate normal distribution parameters μ_0 and Σ_0. Specify smoothing parameter β and q, sample size control parameter α, threshold increase parameter ε, and a continuous and strictly increasing function U (·). Set k=0.
2. Repeat until a specified stopping rule is satisfied:

2.1 Generate N_k i.i.d. samples $(S_1^*)^k,...,(S_{N_k}^*)^k$ from the $N(\hat{\mu}_k, \hat{\Sigma}_k)$ distribution.

2.2 Find the sample $(1 - \rho_k)$-quantile $\gamma_{k+1}(\rho_k, N_k)$ of the samples $\{L(S_i^*)\}^k$, i = 1,…,N_k.

2.3 If k = 0 or $\gamma_{k+1}(\rho_k, N_k) \geq \bar{\gamma}_k + \varepsilon$, then

$$\text{Set } \bar{\gamma}_{k+1} \leftarrow \gamma_{k+1}(\rho_k, N_k), \rho_{k+1} \leftarrow \rho_k, N_{k+1} \leftarrow N_k.$$

Else, find the largest $\bar{\rho} \in (0, \rho_k)$ such that $\gamma_{k+1}(\bar{\rho}, N_k) \geq \bar{\gamma}_k + \varepsilon$.

If such a $\bar{\rho}$ exists, then set $\bar{\gamma}_{k+1} \leftarrow \gamma_{k+1}(\bar{\rho}, N_k), \rho_{k+1} \leftarrow \bar{\rho}, N_{k+1} \leftarrow N_k$.

Else set $\bar{\gamma}_{k+1} \leftarrow \overline{\gamma_k}, \rho_{k+1} \leftarrow \rho_k, N_{k+1} \leftarrow \alpha N_k$

2.4 Update the distribution parameters μ_{k+1} and $\sum_{k+1}$ according to equations (3.4) and (3.5).

2.5 Smooth the parameters by using equation (3.6), (3.7), and (3.8).

2.6 Set k ← k+1.

We choose a strictly increase function U (·), because we consider a maximization problem here. U (·) needs to be a strictly decreasing function for a minimization problem. $\hat{\mu}_k$ and $\hat{\Sigma}_k$ are the parameters after smoothing the μ_{k+1} and $\sum_{k+1}$ originally computed from the samples. In step 2.2, since our goal is a maximization problem, the sample $(1 - \rho_k)$-quantile γ_{k+1} is obtained by first ordering the sample performances $\{L(S_i^*)\}^k$, $i = 1, \ldots, N_k$ from smallest to largest, and then taking the $[(1 - \rho_k)N_k]$th order statistic. Step 2.3 is to find a non-decreasing threshold $\overline{\gamma_k}$, the sample size, and good performance samples selection parameter for the next iteration. A small size of sample paths N_k might cause the algorithm to fail to converge and result in poor quality solutions, while a too large sample size will lead to high computational cost. Large ρ_k tends to use both the "good" and "bad" samples to update the probabilistic model, which slows down the convergence process, whereas a too small ρ might lead to an illogical result. Therefore, we make sample size N_k and threshold proportion parameter ρ_k dynamically adjusted in our algorithm, more specifically, the sample size is adaptively increasing and the parameter is adaptively decreasing as described in step 2.3. α is the rate of sample size increase. A

small positive number ε is selected to ensure that $\{\overline{\gamma}_k\}$ is non-decreasing in the update procedure. At each iteration k, if the new quantile γ_{k+1} is large enough ($\gamma_{k+1}(\rho_k, N_k) \geq \overline{\gamma}_k + \varepsilon$), then we use this quantile as the new threshold and use the current sample size and ρ_k in the next iteration. Otherwise, it indicates that either ρ_k is too large or N_k is too small. First we try to find a smaller $\overline{\rho} < \rho_k$ such that the new sample $(1-\overline{\rho})$ quantile satisfies the above inequality. If such a $\overline{\rho}$ exists, then we decrease the ρ_k value and keep N_k unchanged in the next iteration. If no such $\overline{\rho}$ exists, then we increase the sample size by rate of α while ρ_k and $\overline{\gamma_k}$ remain unchanged. After we find ρ_{k+1}, $N_{k+1,}$ and threshold $\overline{\gamma}_{k+1}$, only those candidate solutions that have better performances than the new threshold will be used in the next iteration.

3.3 Numerical Results

In this section we present numerical results from the MRAS algorithm for both American call/put options and Asian options, respectively. We also give results from the CE method for comparison purpose. All the options in our numerical experiments have a finite number of early exercise opportunities, and are sometimes termed Bermudan derivatives. The stopping criteria at iteration k is 1) cov_max < 1.0, or 2) $\gamma_k = \gamma_{k-1} = \gamma_{k-2}$, or 3) $N_k > N_{max}$. The cov_max is the maximum element in the covariance matrix of the multivariate normal distribution model. It is a measure of the convergence quality. For each test case we use the following parameters: $\rho_0 = 0.5$, N_0

= 100, $\alpha = 2$, $\varepsilon = 10^{-3}$, and $N_{max} = 1000$. Throughout this essay, we assume options are not exercisable at time 0. We choose U (L (S*)) := exp(χL(S*)), where $\chi = 0.1$. The random number generator is taken from Pierre L'Ecuyer's random number package (2002), which offers a better control of streams and seeds. The experiments were implemented with Matlab on a 1.5GHz computer.

3.3.1 American-Style Call Option

The parameters to be estimated in the optimization problem are the critical prices $\{S_i^*\}$, which we obtain by estimating the critical price increments $\{X_i\} \sim N(\mu_k, \Sigma_k)$ at each exercise date, given a starting point S_0^*. After we obtain the optimized increments, the critical prices are computed. Thus, for an option with n exercise dates, we have the following n-1 critical prices to be estimated (the critical price at the last exercise date, the maturity, is known):

$S_1^* = S_0^* + X_1$;

$S_2^* = S_1^* + X_2$;

….

$S_{n-1}^* = S_{n-2}^* + X_{n-1}$.

Therefore, the initial conditions for simulation include the selection of S_0^*, the $\{X_i\}$'s initial mean vector μ_0, and initial variance-covariance matrix $\sum_0$. We set the initial covariance between parameters as 0, and the initial variance is same for all X_i, i.e., $\sum_0$ is a diagonal matrix. The MRAS algorithm is not sensitive to the choice of

initial mean and covariance matrix, provided that the initial sampling variance is chosen large enough. Note that $\sum$ at steps other than the initial one is not necessarily diagonal matrix, because the updating scheme (equation (3.5)) will cause the nonzero covariance between parameters.

We first apply MRAS algorithm to price the American-style call option. We assume the underlying stock price follows geometric Brownian motion:

$$dS_t = (r - \delta)S_t dt + \sigma S_t dW , \tag{3.9}$$

where W_t is a standard Brownian motion process, r is the interest rate, δ is the dividend yield, and σ is the volatility. This leads to the discrete form used in the simulation:

$$S_{t+\Delta t} = S_t \exp((r - \delta - \sigma^2 / 2)\Delta t + \sigma\sqrt{\Delta t}Z , \tag{3.10}$$

where Z is random variate generated from standard normal distribution, i.e., $Z \sim N(0,1)$. Table 3.1 illustrates the price estimates and their 95% confidence intervals based on 1,000,000 replications (50,000 replications each run for 20 different seeds) with obtained parameters of early exercise boundary from simulation, for a 3 year (T = 3) Bermudan call option with r = 0.05, σ = 0.2, δ = 0.04 and K=100. The option is exercisable every 0.5yr (n = 6). We study the performance of MRAS for different initial condition settings: μ_0 = [-5, -5, -5, -5, -5] for initial critical price starting point S_0^* = 130, 140, 150, 160, 170, and 180, [-4, -4, -4, -4, -4] for S_0^* = 120, and [-2, -2, -2, -2, -2] for S_0^* = 110, that are bounded by the lower limit of the critical price at maturity. The diagonal (variance) of $\sum_0$ is 100 for all cases. The circumstances of options considered here include in-the-money (S_0 = 110 and 140), at-the-money (S_0 =

100), and out-of-the-money (S_0 = 60 and 90). Results from MRAS are compared with those from SPSA, PASA and sequential dynamic programming (DP) algorithms presented in Fu et al. (2001). We also compare them with the outcomes from secant and tangent methods described in Laprise et al. (2006); moreover, the corresponding European call option prices are available at the last row of the table.

Our experiments indicate that MRAS algorithm provides an accurate and efficient way to price American call option. It converges to the optimal value within 10 iterations, and the convergence is independent on the initial conditions. We can achieve a cov_max less than 10 in all cases. The results are consistent with the findings of other approaches to similar accuracy, and 95% confidence interval is about 5%. The price of S_0 = 60 is close to European call price, because it is deep out-of-the-money and the possibility of exercise is very small.

Table 3.1 Bermudan Call Option on Asset under Geometric Brownian Motion
K = 100, T = 3.0, n = 6, r = 0.05, δ = 0.04, σ = 0.2

Method	Initial Condition	S0 = 60		S0 = 90		S0 = 100		S0 = 110		S0 = 140	
		Price	C.I.	Price	C.I.	Price	C.I.	Price	C.I.	Price	C.I.
MRAS	110	0.87	0.01	8.64	0.03	13.56	0.04	19.52	0.04	42.32	0.06
	120	0.87	0.01	8.64	0.03	13.56	0.04	19.52	0.04	42.32	0.06
	130	0.86	0.01	8.65	0.03	13.57	0.03	19.52	0.04	42.33	0.06
	140	0.87	0.02	8.64	0.03	13.56	0.03	19.52	0.04	42.28	0.09
	150	0.88	0.02	8.64	0.04	13.57	0.04	19.52	0.04	42.33	0.06
	160	0.87	0.01	8.65	0.03	13.56	0.04	19.51	0.04	42.33	0.06
	170	0.87	0.01	8.64	0.03	13.56	0.04	19.52	0.04	42.32	0.06
	180	0.87	0.01	8.65	0.03	13.57	0.04	19.52	0.04	42.32	0.06
SPSA						13.69	0.04				
PASA						13.65	0.04				
DP						13.39	0.04				
Secant		0.88		8.63		13.56		19.53		42.29	
Tangent		0.87		8.63		13.55		19.53		42.29	
Eur		0.87		8.55		13.37		19.18		40.74	

Table 3.2 displays the thresholds for $S_0 = K = 100$ at t = 0.5, 1.0, 1.5, 2.0, and 2.5. The deviation between the obtained optimal prices related to the initial settings is relatively small, specifically, values fluctuating within the range of less than 5%. It is important to note that this table illustrates an at-the-money American call option example, where the fluctuation is expected to be large for the property of at-the-money option itself. Our results from other scenarios suggest an even smaller critical price region depended on S_0*.

Table 3.2 Thresholds of Bermudan Call Option
$K = S_0 = 100$, T = 3.0, n = 6, r = 0.05, $\delta = 0.04$, $\sigma = 0.2$

Method	Initial Condition	t = 0.5	t = 1.0	t = 1.5	t = 2.0	t = 2.5
MRAS	110	155.05	153.74	151.36	148.20	140.69
	120	153.35	151.69	148.50	146.70	133.22
	130	158.55	155.40	150.46	144.04	132.96
	140	157.06	151.27	148.67	143.94	130.80
	150	153.77	152.64	149.59	144.47	126.46
	160	158.04	154.36	147.74	145.58	129.43
	170	162.47	156.70	152.65	148.91	136.35
	180	157.11	150.56	147.60	144.28	132.88
Secant		158.43	154.06	148.68	141.70	
Tangent		158.42	154.05	148.67	141.70	

3.3.2 American-Style Put Option

3.3.2.1. Underlying Asset Follows Geometric Brownian Motion

We now consider the case of implementing MRAS algorithm for pricing American-style put option whose underlying stock price follows geometric Brownian motion (eq.(3.10)). The optimized parameter setting of American put option is similar to that of American call option.

Shown in Table 3.3 is the price of a 3 year (T = 3) American put options with $r = 0.05$, $\sigma = 0.2$, $\delta = 0$, $K = 100$, and $n = 6$. μ_0 is [5, 5, 5, 5, 5] for $S_0^* = 30$, 40, 50, 60, and 70, [4, 4, 4, 4, 4] for $S_0^* = 80$, and [2, 2, 2, 2, 2] for $S_0^* = 30$, according to the upper limit of the critical price at maturity. The diagonal (variance) of $\sum_0$ is 100. Like the example of American call option, various scenarios of in-the-money ($S_0 = 60$ and 90), at-the-money ($S_0 = 100$), and out-of-the-money ($S_0 = 100$ and 140) are examined. The results from secant and tangent methods of Laprise (2006) are listed for comparison. Analogous to Table 3.2, Table 3.4 presents the threshold estimates for at-the-money put option for each choice of S_0^*.

Our experiments find consistently the maximum values regardless of the initial choices, indicating that the true global optimum is reached in each case. For some initial settings, we add constraint $\rho > 0.1$ so that the algorithm can avoid a local optimum or an extreme case from stochastic simulation. The algorithm approaches to the optimal value within 15 iterations for most cases, and less than 5 for $S_0 \neq K$ with an initial critical price close to the optimum ($S_0^* = 50$, 60, 70, 80, 90). The at-the-money prices are the least accurate, e.g., for $S_0 = 110$, only 5 iterations are needed for convergence, while more than 10 iterations are required for comparable accuracy for $S_0 = 100$. We also find the threshold bounds for the put options are tighter than for the calls (Table 3.2) as shown in Table 3.4.

Table 3.3 Bermudan Put Option on Asset under Geometric Brownian Motion
K = 100, T = 3.0, n = 6, r = 0.05, $\delta = 0$, $\sigma = 0.2$

Method	Initial Condition	S0 = 60	S0 = 90	S0 = 100	S0 = 110	S0 = 140
MRAS	30	37.31	12.96	8.39	5.52	1.51
	40	37.48	12.95	8.39	5.52	1.48
	50	37.51	12.91	8.40	5.52	1.52
	60	37.52	12.95	8.39	5.52	1.51
	70	37.52	12.94	8.39	5.50	1.51
	80	37.52	12.95	8.43	5.52	1.51
	90	37.53	12.97	8.45	5.49	1.53
Secant		37.55	12.91	8.45	5.50	1.50
Tangent		37.55	12.91	8.45	5.50	1.50
European		27.97	10.24	7.00	4.71	1.37

Table 3.4 Thresholds of Bermudan Put Option
$K = S_0 = 100$, T = 3.0, n = 6, r = 0.05, $\delta = 0$, $\sigma = 0.2$

Method	Initial Condition	t = 0.5	t = 1.0	t = 1.5	t = 2.0	t = 2.5
MRAS	30	81.87	84.29	86.63	88.54	90.64
	40	81.50	84.38	87.62	88.86	90.22
	50	83.63	85.27	85.87	86.89	89.02
	60	80.30	83.38	84.78	86.07	89.61
	70	83.86	85.31	87.83	88.37	90.33
	80	80.57	81.99	84.49	88.07	89.19
	90	81.60	82.29	85.06	86.65	89.00
Secant		83.06	84.02	85.32	87.20	
Tangent		83.06	84.3	85.32	87.20	

3.3.2.2. Underlying Asset Follows Merton Jump Diffusion Model

The jump-diffusion process is appealing because it allows price discontinuities and addresses the issue of 'fat tails', but the presence of random jumps complicates the valuation of the American put option. As a Monte Carlo simulation method, MRAS

algorithm is capable of effectively incorporating the jump process. The Merton jump diffusion model is formulated as:

$$\frac{dS_t}{S_t} = (r - \delta)dt + \sigma dW + J_t dq$$

where dq is a Poisson random variable that takes value zero (no jump realized) with probability 1-λdt, and value one (jump realized) with probability λdt, and J is a stochastic jump size.

The discrete form used in simulation can be written as follows:

$$S_{t+\Delta t} = S_t \exp((r - \delta - \sigma^2/2)\Delta t + \sigma\sqrt{\Delta t}Z_0 + \sum_{j=1}^{N(\Delta t)}(\gamma Z_j - \gamma^2/2)) \qquad (3.11)$$

where $Z_j \sim N(0,1)$ i.i.d., $N(\Delta t) \sim$ Poisson $(\lambda\Delta t)$ is the number of jumps within time Δt, the jump sizes are i.i.d. lognormally distributed: $LN(-\gamma^2/2, \gamma^2)$, λ is the jump frequency, and γ is the jump volatility.

Table 3.5 shows the results of applying the MRAS algorithm to a six-month (T = 0.5yr) put option written on a single stock modeled by the jump-diffusion model without dividend ($\delta = 0$), and r = 0.1, $\sigma = 0.2828$, $\lambda = 2$, $\gamma = 0.2$, S0 = K = 100. European price (n = 1) for this example is 8.393. After we obtained the early exercise thresholds, we simulate the pricing process by 50,000 replications each run for 20 different seeds, and 95% confidence interval is calculated and the value of the confident interval is within 0.02 - 0.04.

We compare the outcome of MRAS with other algorithms including SPSA, DP, and Secant/Tangent interpolation methods, and we found the MRAS results are closest to those of secant/tangent algorithms, and the values are between those from SPSA and DP when n is small (n = 2, 3) and they are more consistent as n increases.

Moreover, Secant method provides the upper bound for the results. MRAS is an efficient way to price American put option written on jump-diffusion process since the simulations converge to the optimal value within 20 iterations regardless of the initial choice of S_0*.

Table 3.5 Bermudan Put Option on Asset under Merton Jump-Diffusion
$S_0 = 100$, $K = 100$, $T = 0.5$, $r = 0.1$, $\delta = 0$, $\sigma = 0.2828$, $\lambda = 2$, $\gamma = 0.2$

Method	Initial Condition	n = 2		n = 3		n = 4		n = 6	
		Initial	Final	Initial	Final	Initial	Final	Initial	final
MRAS	30	8.38	8.56	8.37	8.65	8.32	8.62	8.39	8.73
	40	8.37	8.56	8.38	8.65	8.34	8.63	8.43	8.72
	50	8.39	8.57	8.41	8.64	8.38	8.63	8.53	8.73
	60	8.42	8.57	8.48	8.66	8.49	8.63	8.68	8.71
	70	8.51	8.57	8.61	8.63	8.61	8.63	8.64	8.73
	80	8.59	8.58	8.62	8.65	8.46	8.62	8.22	8.71
	90	8.27	8.58	7.75	8.65	8.04	8.63	7.58	8.73
SPSA			8.49		8.62		8.70		
DP			8.57		8.88		8.73		
Secant			8.61		8.69		8.73		8.77
Tangent			8.61		8.68		8.72		8.76

Table 3.6 shows the threshold range estimated on different S_0* for the American put option written on jump-diffusion model with various number of exercise opportunities. For n = 6, the critical prices are similar to those without jumps. The algorithm converges efficiently because the upper and low limit of the range is relatively small, less than 5% of the mean critical price, independent of the initial setting.

Table 3.6 Thresholds of Bermudan Put Option on Asset under Merton Jump-Diffusion $S_0 = 100$, $K = 100$, $T = 0.5$, $r = 0.1$, $\delta = 0$, $\sigma = 0.2828$, $\lambda = 2$, $\gamma = 0.2$

n					
2	t = 1/4				
	80.0 ~ 83.3				
3	t = 1/6	t = 1/3			
	77.7 ~ 82.9	83.5 ~ 85.6			
4	t = 1/8	t = 1/4	t = 3/8		
	78.3 ~ 81.1	81.0 ~ 83.7	84.3 ~ 85.9		
6	t = 1/12	t =1/6	t = 1/4	t = 1/3	t = 5/12
	77.7 ~ 79.9	80.5 ~ 82.2	82.2 ~ 84.5	82.7 ~ 88.7	86.7 ~ 90.5

3.3.3 American-style Asian Call Option

Trading in Asian options grows rapidly because Asian options provide payoffs with average property that may well match risk management characteristics. Hull and White (1993) applied modified binomial and trinomial lattices to value American-style Asian options. Their method is limited in its applicability since only an average beginning at time zero can be handled. Due to the dependence on the entire path of the underlying asset, no explicit formula for the distribution of the average price exists yet and Asian options appear to be particularly suited for Monte Carlo simulation. Grant et al. (1997) incorporated optimal early exercise in the Monte Carlo method of valuing American-style Asian options by linking forward-moving simulation and the backward-moving recursion of dynamic programming. They identified the locus of critical prices for American-Asian options by equating the holding value and early exercise payoff from a preset finite parameter grid. Following a piecewise linear approximation of the exercise boundary illustrated in Grant et al.

(1997), Fu et al. (2001) applied stochastic approximation with gradient estimators to characterize the early exercise thresholds for American-Asian options.

Asian options have payoffs that depend on the average value of the underlying asset at some specified time point. The payoff function at early exercise date t_i of American-Asian call option is defined as

$$L_i = (\bar{S}_i - K)^+$$

where $\bar{S}_i = \frac{1}{t_i - t_0 + 1} \sum_{j=t_0}^{t_i} S_j$ is the arithmetic average of the stock prices from date t_0 up to the exercise time t_i. For the American-Asian call option, we can still employ the maximize payoff function (equation (3.1)) with $\bar{S}_i$ instead of S_i, and the option should be exercised when the average asset price is greater than the exercise critical price.

Different from American call and put options, the early exercise boundary of Asian options relies on both the average asset price and the current asset price. There is a critical average price for each asset price S_i, and the early exercise thresholds are expressed as a locus of critical prices, $\bar{S}_i^*(S_i)$. Two parameters are used to estimate critical price $\bar{S}_i^*$ at each exercise point in Grant et al. (1997) and Fu et al. (2001). We use only one parameter to estimate the locus of critical prices as

$\bar{S}_i^* = S_i + X_i$, if $S_i > K$;

$\bar{S}_i^* = K + X_i$, if $S_i \le K$.

It's a one-parameter linear approximation of the early exercise boundary; however, we will show later that the simulation results are reasonable with only one parameter and comparable with the results from two-parameter model. This demonstrates the conclusion from Grant et al. (1997) that the estimated value of the American-style Asian call option is relatively insensitive to modest errors in the early exercise boundaries.

We followed numerical examples from Grant et al. (1997) and Fu et al. (2001), where the early exercise opportunities are the discrete points $\{t_j, j = 1, \ldots, N\}$. More specifically, we have maturity T = 120 days, averaging starts at day 91 ($t_0 = 91$), and the earliest exercise is day 105. We test three settings of early exercise opportunities: {105, 120}, {105, 110, 115, 120}, and {105, 108, 111, 114, 117}. In addition, we use the initial stock price $S_0 = 100$, strike price K = 90, 95, 100, 105, 110, interest rate r = 0.09, dividend $\delta = 0$, and volatility $\sigma = 0.2$ or 0.3. We assume the underlying stock price process follows geometric Brownian motion according to equation (3.10). We run the simulation based on 100,000 generated samples. Results from MRAS method are compared with those from SPSA, PASA and DP methods. Table 3.7 shows American-Asian option prices for varying strike prices and exercises opportunities as well as the 95% confidence interval when $\sigma = 0.2$, and Table 3.8 shows the results for $\sigma = 0.3$. We find that the MRAS algorithm converges very quickly by obtaining the optimal parameters within 10 iterations for most cases.

Table 3.7 American-Asian Call Option on Asset under Geometric Brownian Motion
$S_0 = 100$, $T = 120$, $r = 0.09$, $\delta = 0$, $\sigma = 0.2$

	t_i = 105, 120		t_i = 105, 110,115,120		t_i=105, 108,111,…,120	
	Price	95% C. I.	Price	95% C. I.	Price	95% C. I.
K = 90						
MRAS	**13.06**	**0.03**	**13.12**	**0.04**	**13.10**	**0.03**
SPSA	13.11	0.02	13.04	0.02	13.17	0.02
PASA	13.09	0.02	13.18	0.02	13.20	0.02
DP	13.08	0.02	13.17	0.02	13.19	0.02
K = 95						
MRAS	**9.01**	**0.03**	**9.04**	**0.03**	**9.10**	**0.03**
SPSA	8.98	0.02	8.97	0.02	9.05	0.02
PASA	9.02	0.02	9.11	0.02	9.12	0.02
DP	9.02	0.02	9.10	0.02	9.12	0.02
K = 100						
MRAS	**5.70**	**0.02**	**5.71**	**0.02**	**5.71**	**0.02**
SPSA	5.67	0.01	5.74	0.01	5.69	0.01
PASA	5.71	0.01	5.77	0.01	5.79	0.01
DP	5.70	0.01	5.77	0.01	5.79	0.01
K = 105						
MRAS	**3.28**	**0.02**	**3.30**	**0.02**	**3.32**	**0.02**
SPSA	3.22	0.01	3.28	0.01	3.20	0.01
PASA	3.29	0.01	3.33	0.01	3.34	0.01
DP	3.28	0.01	3.33	0.01	3.34	0.01
K = 110						
MRAS	**1.68**	**0.01**	**1.73**	**0.01**	**1.72**	**0.01**
SPSA	1.64	0.01	1.66	0.01	1.66	0.01
PASA	1.72	0.01	1.75	0.01	1.75	0.01
DP	1.72	0.01	1.75	0.01	1.75	0.01

Table 3.8 American-Asian Call Option on Asset under Geometric Brownian Motion
$S_0 = 100$, $T = 120$, $r = 0.09$, $\delta = 0$, $\sigma = 0.3$

	t_i = 105, 120		t_i = 105, 110,115,120		t_i=105, 108,111,…,120	
	Price	95% C. I.	Price	95% C. I.	Price	95% C. I.
K = 90						
MRAS	**14.35**	**0.05**	**14.42**	**0.05**	**14.46**	**0.05**
SPSA	14.20	0.03	14.23	0.03	14.25	0.03
PASA	14.38	0.03	14.50	0.03	14.53	0.03
DP	14.37	0.03	14.49	0.03	14.53	0.03
K = 95						
MRAS	**10.78**	**0.04**	**10.86**	**0.04**	**10.91**	**0.04**
SPSA	10.65	0.02	10.73	0.03	10.83	0.03
PASA	10.82	0.03	10.91	0.03	10.94	0.03
DP	10.80	0.03	10.91	0.03	10.94	0.03
K = 100						
MRAS	**7.80**	**0.03**	**7.87**	**0.03**	**7.87**	**0.03**
SPSA	7.68	0.02	7.93	0.02	7.81	0.02
PASA	7.82	0.02	7.92	0.02	7.93	0.02
DP	7.81	0.02	7.92	0.02	7.94	0.02
K = 105						
MRAS	**5.44**	**0.02**	**5.49**	**0.03**	**5.50**	**0.03**
SPSA	5.46	0.02	5.51	0.02	5.44	0.02
PASA	5.45	0.02	5.53	0.02	5.54	0.02
DP	5.45	0.02	5.53	0.02	5.54	0.02
K = 110						
MRAS	**3.66**	**0.02**	**3.69**	**0.02**	**3.70**	**0.02**
SPSA	3.61	0.02	3.67	0.02	3.59	0.02
PASA	3.67	0.02	3.72	0.02	3.73	0.02
DP	3.66	0.02	3.73	0.02	3.74	0.02

3.3.4 Comparison between MRAS and CE Methods

Both MRAS and CE are model-based methods, which start with a parameterized probability distribution on the solution space and update the parameters at each iteration towards a 'better' solution. KL-divergence is used as a measure and is expected to be minimized. In MRAS, a sequence of reference distributions is adopted and the minimum KL-divergence is achieved between the next step distribution and

the current reference model, while in CE a single optimal distribution is used and KL-divergence measures the distance between the optimal distribution and the family of parameterized distributions.

The CE algorithm works as follows:

Algorithm CE: Continuous Optimization – Monte Carlo version

---

1. Initialize: Specify quantile parameter ρ and sample size N. Initialize parameters of the probabilistic model (multivariate normal distribution) μ_0 and Σ_0. Set k=0.
2. Repeat until a specified stopping rule is satisfied:
    a. Generate N i.i.d. samples $X_1, \ldots, X_N$ from the $N(\hat{\mu}_k, \hat{\Sigma}_k)$ distribution.
    b. Select the ρN best performing (elite) samples, and let I be the indices of the ρN best performing samples.
    c. Update the parameters as:

$$\mu_{k+1} = \frac{1}{\rho N} \sum_{i \in I} X_i \text{ , and}$$

$$\Sigma_{k+1} = \frac{1}{\rho N} \sum_{i \in I} (X_i - \mu_{k+1})(X_i - \mu_{k+1})^T$$

    d. Smooth by using equation (3.6), (3.7) and (3.8).
    e. Set k ← k+1.

In CE method, a fixed number (ρN) of best performing samples is selected at each iteration, while the quantile parameter ρ_k and sample size N_k keep updating in MRAS. ρ_k in MRAS is a proportion of samples that will be used to update the

probabilistic model, and the purpose is to concentrate the computational effort on the set of elite samples. We study the sensitivity on choice of initial ρ using an example of Bermudan put option written on a single asset following geometric Brownian motion, and the model setting is $K = 100$, $T = 3.0$, $N = 6$, $r = 0.05$, $\delta = 0$, $\sigma = 0.2$, the initial critical prices [35, 40, 45, 50, 55], and initial covariance matrix with 100 on the diagonal and 0 otherwise. ρ remains unchanged through the CE simulation and varies according to some updating mechanism in MRAS.

Figures 3.1 and 3.2 display the thresholds evolution from CE and MRAS when initial $\rho = 0.2$, and we found CE method gives a smoother convergence while MRAS converges a little faster. Both algorithms provide sound results efficiently. The graphical representations of the critical prices convergence of CE and MRAS in the case of $\rho = 0.5$ are given in Figures 3.3 and 3.4, respectively. It is not difficult to see that MRAS approaches to the optimal value much quicker than CE method, more specifically, MRAS reaches the optimum at iteration 5 whereas it takes 20 iterations for CE. Figure 3.4 also shows the sequence of $(1- \rho_k)$ quantile of MRAS, and it increases as expected since this is a maximization problem. As mentioned before, the maximum element in the covariance matrix of the multivariate normal distribution model (cov_max) is one measure of the convergence. Figure 3.5 illustrates the progress of cov_max from MRAS and CE algorithms when $\rho = 0.5$. It is not surprise to see that cov_max of MRAS declines much faster though fluctuates more, e.g., it reaches 450 for one instance, while those for CE are less than 250. Figures 3.7 and 3.8 present the similar thresholds evolution for $\rho = 0.8$. It is obvious that CE method doesn't converge well, but MRAS reaches the optimum efficiently despite the large ρ

we choose. The critical prices from CE don't show convergence after even 200 iterations while MRAS satisfy terminate conditions within 30 steps. Figure 3.6 compares the cov_max from MRAS with CE method for $\rho = 0.8$. The left y axis is for MRAS and right y-axis is for CE. The graph implies a similar finding as the one suggested by the example of $\rho = 0.5$, but in a more extreme instance: MRAS converges significantly better and efficient than CE and the development of cov_max is flatter in CE algorithm. Tables 3.9 and 3.10 present the evolution of MRAS and CE algorithms for $\rho = 0.8$, respectively. For MRAS, ρ_k decreases from the initial 0.8 to 0.11 at termination point, cov_max is reduced by 2 magnitudes, and number of sample size doubles at the last iteration, whereas the cov_max from CE keeps relatively smooth with most in the range of 70 to 200, and the critical prices moves slowly.

From the analysis, we can conclude that ρ assumes an important role in the optimization process of CE method. Unlike MRAS, where the convergence of the sequence of reference models to an optimal distribution model is guaranteed, the convergence of the sequence in CE relies on the quantile parameter ρ. CE method can obtain a favorable result only when the value of ρ is chosen sufficiently small because an importance sampling technique is simply and solely employed in the parameter updating procedure. In contrast, the MRAS algorithm is insensitive to the choice of initial quantile parameter and sample size, since both parameters will adapt corresponding to the updating schemes in the successive iterations.

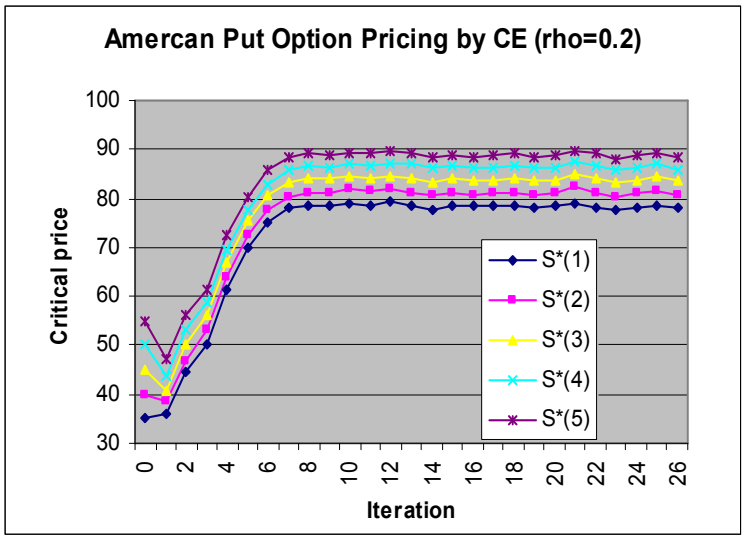

Fig.3.1.Critical price of CE (ρ=0.2)

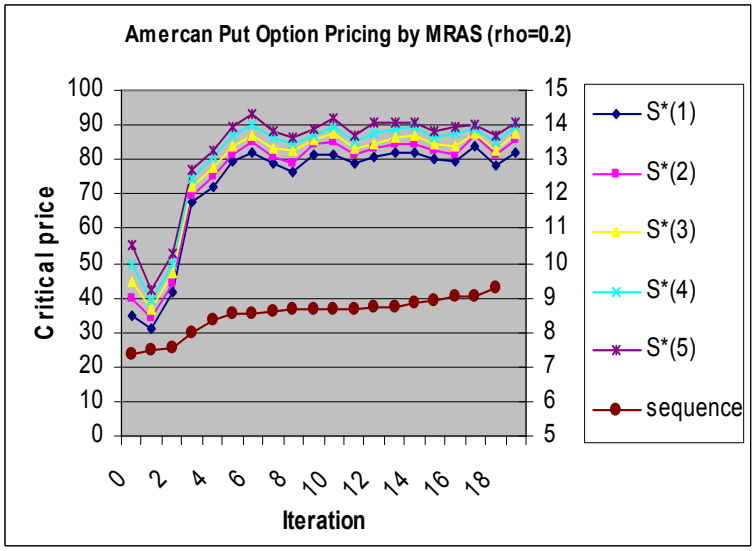

Fig.3.2.Critical price and quantile of MRAS (ρ=0.2)

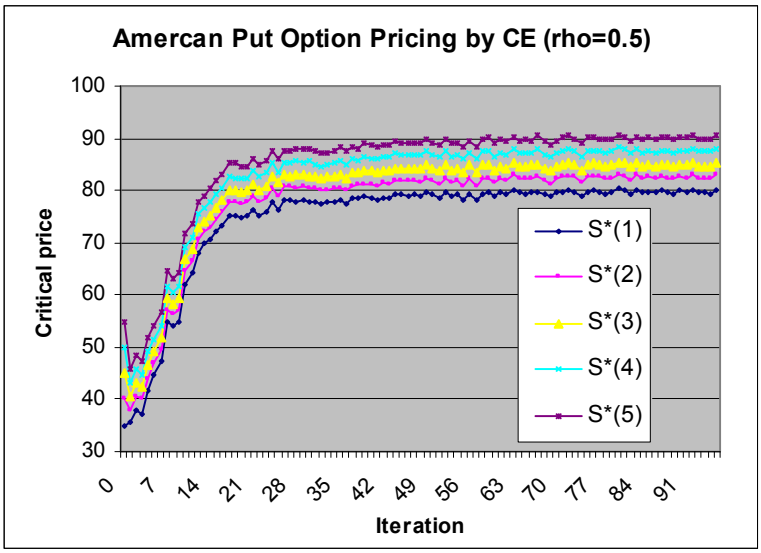

Fig.3.3. Critical price of CE (ρ=0.5)

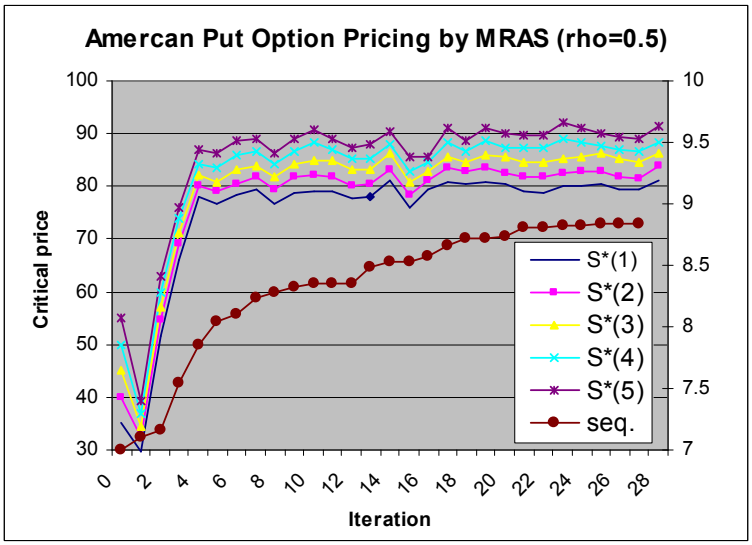

Fig.3.4. Critical price and quantile of MRAS (ρ=0.5)

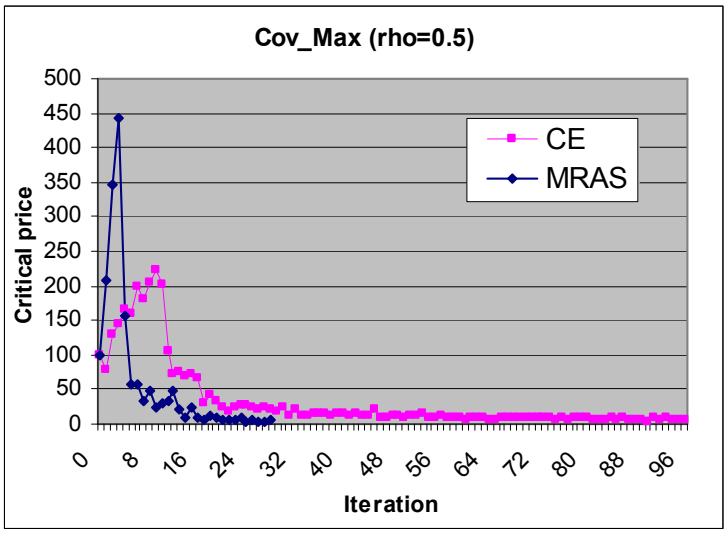

Fig.3.5.Cov_max of CE and MRAS (ρ=0.5)

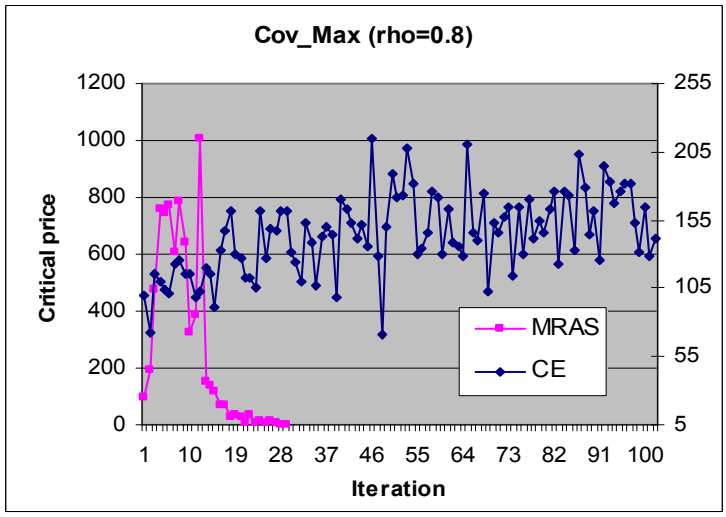

Fig.3.6.Cov_max of CE and MRAS (ρ=0.8)

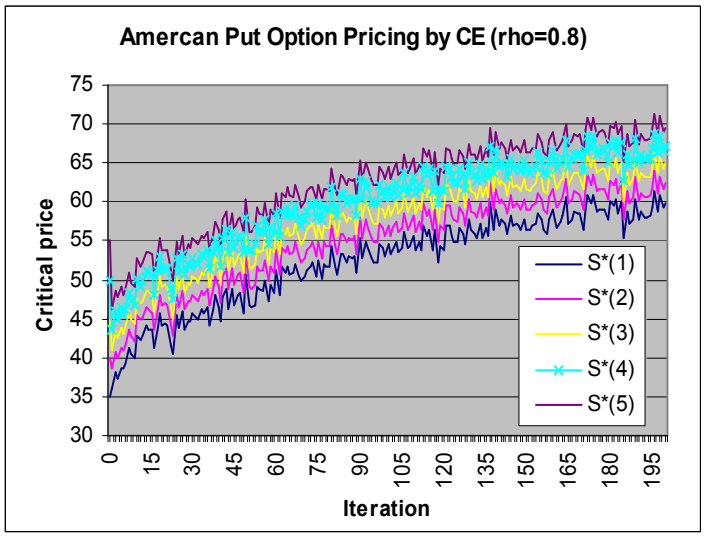

Fig.3.7.Critical Price of CE (ρ=0.8)

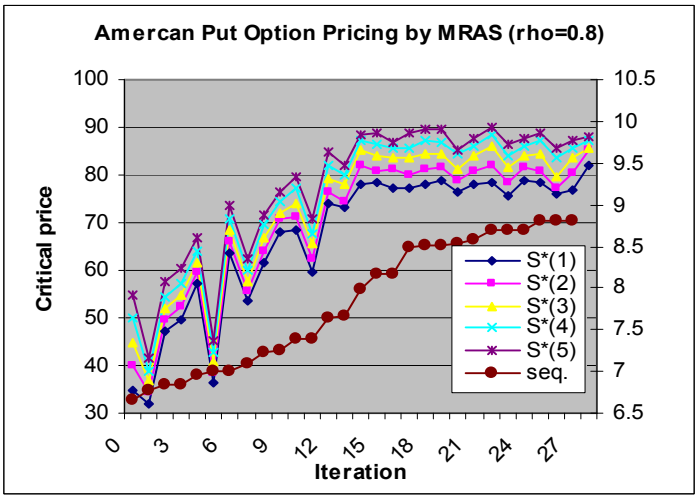

Fig.3.8.Critical Price and quantile of MRAS (ρ=0.8)

Table 3.9 The evolution of MRAS algorithm for $\rho_0 = 0.8$

iter k	rho	N	quantile	cov_max	Critical Price				
0				100.00	35.00	40.00	45.00	50.00	55.00
1	0.8	100	6.66	193.99	31.94	34.49	36.83	38.95	41.44
2	0.8	100	6.78	477.48	47.17	49.75	52.18	54.60	57.76
3	0.8	100	6.85	756.34	49.63	52.38	54.82	57.38	60.20
4	0.8	100	6.85	743.62	57.36	59.57	61.74	64.15	66.89
5	0.76	100	6.95	774.53	36.55	38.41	40.95	42.93	45.06
6	0.76	100	6.99	607.75	63.76	66.03	68.45	70.84	73.52
7	0.76	100	6.99	782.79	53.43	55.57	57.73	60.19	62.48
8	0.67	100	7.10	638.92	61.55	64.04	66.67	69.41	71.75
9	0.67	100	7.24	322.54	67.88	70.61	72.18	74.41	76.57
10	0.67	100	7.26	386.77	68.52	71.24	74.05	77.21	79.79
11	0.62	100	7.38	1004.31	59.60	62.40	65.63	67.68	70.64
12	0.62	100	7.39	150.27	74.04	76.37	79.36	82.11	84.85
13	0.48	100	7.64	134.87	73.22	74.47	77.84	79.99	81.89
14	0.48	100	7.67	116.94	77.93	81.93	85.39	87.04	88.48
15	0.31	100	7.98	70.53	78.49	80.75	83.84	86.47	88.71
16	0.31	100	8.17	66.38	77.32	81.04	83.66	85.49	86.92
17	0.31	100	8.18	28.78	77.34	79.92	83.42	85.56	88.76
18	0.16	100	8.49	33.71	78.20	81.20	84.31	87.02	89.80
19	0.16	100	8.51	24.66	78.82	81.61	84.30	86.98	89.65
20	0.15	100	8.51	10.33	76.29	78.89	81.39	84.29	85.19
21	0.13	100	8.53	32.02	78.10	80.72	83.87	85.99	87.77
22	0.12	100	8.59	4.55	78.23	82.13	85.93	88.52	89.95
23	0.12	100	8.69	11.47	75.50	78.53	81.49	84.01	86.31
24	0.12	100	8.69	4.36	78.64	81.61	84.10	85.84	87.72
25	0.11	100	8.70	13.32	78.22	80.98	84.32	87.29	88.87
26	0.11	100	8.81	4.61	75.86	77.15	79.41	83.58	85.48
27	0.11	100	8.81	2.19	76.85	80.36	83.72	85.46	87.09
28	0.11	200	8.81	2.75	82.08	85.19	85.73	87.03	88.09

Table 3.10 The evolution of CE algorithm for $\rho = 0.8$

iter k	cov_max	Critical Price				
0	100.00	35.00	40.00	45.00	50.00	55.00
1	72.60	36.03	38.63	41.12	43.67	46.23
5	100.69	38.60	40.88	43.07	45.43	48.03
10	98.96	42.84	45.33	47.59	50.12	52.72
15	132.54	43.58	45.82	48.13	50.57	53.21
20	113.32	44.43	46.99	49.25	51.57	53.87
25	149.23	43.88	46.41	48.71	50.99	53.36
30	123.54	45.75	48.18	50.41	52.88	55.38
35	143.50	46.40	48.93	51.35	53.71	56.13
40	163.27	44.54	47.06	49.46	51.83	54.24
45	214.35	46.87	49.41	51.77	54.02	56.44
50	171.89	46.64	49.20	51.64	53.87	56.19
55	134.90	48.71	51.26	53.70	56.25	58.38
60	163.51	51.27	53.86	56.26	58.59	61.02
65	145.24	50.68	53.32	55.73	58.24	60.87
70	145.20	50.30	52.85	55.35	57.75	60.12
75	130.16	52.43	54.96	57.33	59.73	62.14
80	162.61	54.37	56.84	59.36	62.01	64.31
85	133.25	52.16	54.82	57.42	59.85	62.31
90	125.56	55.32	57.96	60.48	62.89	65.21
95	181.41	52.11	54.74	57.20	59.77	62.15
100	129.14	54.53	57.05	59.45	61.89	64.41
105	133.59	53.98	56.65	59.18	61.74	64.11
110	132.67	53.86	56.26	58.64	61.02	63.42
115	183.10	56.50	59.14	61.67	64.21	66.59
120	117.98	53.83	56.35	58.99	61.37	63.78
125	147.28	54.88	57.48	59.99	62.55	64.98
130	137.97	56.37	58.89	61.31	63.78	66.17
135	116.85	56.71	59.34	61.81	64.22	66.84
140	104.32	57.04	59.76	62.19	64.50	66.92
145	124.33	56.91	59.41	61.97	64.31	66.72
150	134.26	56.48	59.17	61.63	64.01	66.33
155	133.75	57.67	60.34	62.90	65.34	67.76
160	120.70	57.78	60.29	62.79	65.27	67.62
165	111.36	57.16	59.78	62.19	64.66	67.07
170	134.84	56.52	59.05	61.55	64.11	66.47
175	93.82	58.84	61.41	63.79	66.08	68.44
180	88.70	59.98	62.62	65.07	67.57	69.83
185	108.25	55.46	57.86	60.12	62.72	65.21
190	83.37	57.62	60.20	62.71	65.16	67.52
195	94.05	59.35	61.99	64.49	66.96	69.43
200	110.25	59.93	62.47	64.96	67.28	69.57

3.4 Conclusions

This essay applies a stochastic approach, MRAS algorithm, for the pricing of American style options. The method casts the pricing problem as an optimization problem, and optimizes the early exercise thresholds simultaneously by iterative updates via a reference model. We study the case of American call option, American put option and Asian option, which are written on the underlying assets following geometric Brownian motion or jump-diffusion processes. In repeated experiments the global maximum is consistently found for varying initial condition settings. We demonstrate its accuracy and efficiency and give an example where MRAS provides a better solution than CE method. We focus on American-style options with a single underlying asset in this essay, but this methodology can be applied to other types of options, especially the derivatives with complicated exercise regions or higher dimensional options. We can conclude that MRAS is a flexible and useful randomized optimization algorithm in pricing derivatives.

Chapter 4

Optimizing Importance Sampling Parameter for Portfolios of Credit Risky Assets

This essay is organized in the following manner. Section 1 introduces the portfolio model, the application of Monte Carlo simulation and importance sampling technique, and the optimization problem setting. We then describe two stochastic methods used to find optimal importance sampling measure change in Section 2. The derivation of the gradient estimator and the detailed implementation of the approaches are also illustrated in this section. Section 3 sets up the numerical experiment and presents the numerical results of various descriptive statistics outputs. We also compare the performance of IPA and CE methods. Section 4 concludes the essay with a summary and a direction for further study.

4.1 Portfolio Model and Problem Setting

4.1.1 Portfolio Model

Quantitative methods for portfolio analysis have been developed since Markowitz's (1952) pioneering work in 1950s, and have been applied effectively in a variety of areas of finance, particularly to equity portfolios. Similar progress has not occurred for debt portfolios because of the difficulty to quantify the level of default risk in a single asset, and to identify the correlation between the various default risks. Merton (1974) made use of the Black and Scholes (1973) option pricing model to model a single company's credit risk by characterizing the company's equity as a call option on its assets and assuming the firm value follows a random process similar to the one describes generic stocks in equity markets. Kealhofer and Bohn (2001) developed a credit risk model in this framework to manage portfolio of default risk. In the paper, they described methods to measure probabilities of default for each asset, the recovery in the event of default, and the default relationship between the assets in the portfolio. We adopt this model to study the optimal choice of the importance sampling measure parameter in this essay.

The basic idea in the default model is that firms default when their asset return falls below a certain threshold over a fixed time horizon, where the firm asset returns are drawn from a multivariate normal distribution with a given correlation matrix. The loss associated with a default depends on the recovery rate, which refers to the fraction of the amount may be recovered through bankruptcy proceedings or some other form of settlement. Loss given default (LGD) is typically expressed as a proportion of the nominal, and is equivalent to one minus recovery rate. In our model,

we assume the LGD is known and follows an independent and identically Beta distribution, because Beta distribution can be bound between two points and can assume a wide range of shapes.

Default correlation measures the strength of the default relationship between two assets, and can be estimated via historical data of firms' equity return correlations. We turn to a factor model to calculate the correlations, and the firm asset returns e_i* is expressed as

$$e_i^* = \sqrt{r_i^2}\,\beta_i^T z + \sqrt{1-r_i^2}\,\varepsilon_i, \tag{4.1}$$

where

z: a vector of the systematic risks with M elements. The component could represent various risks such as global economic, regional, sector, country, and industry factors.

r_i^2 : the firm's R-square value, that is, the percentage of the asset return variance explained by the total systematic risks.

β_i: a vector of the normalized weights of the systematic risks, and $\beta_i = (\beta_1^i, \beta_2^i, ..., \beta_M^i)^T$.

ε_i: the firm-specified risk.

We build a Monte Carlo simulation that draws the inputs in equation (4.1) repeatedly to determine the portfolio loss. If we define Γ to be the diagonal matrix with $\Gamma_{ii} = r_i^2$, let B to be the matrix B = [β_1, …, β_M], and I as the identity matrix, then we can rewrite equation (4.1) as

$$e^* = \Gamma^{1/2} B^T z + (I-\Gamma)^{1/2}\varepsilon, \tag{4.2}$$

and the correlation matrix becomes

$$P = \Gamma^{1/2} B^T B \Gamma^{1/2} + I - \Gamma . \tag{4.3}$$

A firm defaults when its asset value falls below a certain default threshold, and the loss to the firm is LGD. Thus, the loss of the portfolio with N risky assets is given by

$$L = \frac{1}{N} \sum_{i=1}^{N} LGD_i I\{(\alpha_i - e_i^*) \geq 0\}, \tag{4.4}$$

where

N: number of risky assets in the portfolio,

LGD_i: loss given default of asset i,

I{.}: indicator function,

α_i: default threshold that is related to the default probability p_i by the standard cumulative normal distribution function.

e_i^*: firm asset returns that are drawn from a multivariate normal distribution with a given correlation matrix.

Some assumptions in this portfolio model include the following:

1. We implement a simulation over a single time horizon and assume the maturity of all the instruments in the portfolio is the same as this horizon.
2. The portfolio is homogeneous in exposure size, i.e., each asset accounts for 1/N of the total exposure for a portfolio of N credit risky assets.
3. All the credit exposures are priced at par, so that no uncertain discount factor needs to be considered in the model.
4. There is no correlation between LGD and the asset return for a defaulted loan.

5. All the credit exposures are issued by distinct firms and the firm specified risks are independent.

In this essay, we will study both the expected loss and unexpected loss for the credit risk portfolio. The expected loss (EL) is simply the expected value of portfolio losses due to default over a certain horizon, whereas the unexpected loss (UL) measures the second moment of portfolio losses and it can be either the volatility or some measure of the loss distribution tail such as quantile of portfolio loss. In this essay, the UL refers particularly to the volatility of loss. By the definition, the EL and UL are calculated as

$$E(L) = \int L(e^*) f(e^*) de^* ,$$

and

$$UL = E[(L - EL)^2] = E(L^2) - E^2(L) .$$

In an attempt to measure the loss at a specified probability level, we also estimate the quantile of the loss distribution. The calculation of the loss quantile gives financial institutions the market risk of the portfolio, and it can be used to calculate Value-at-risk that is widely used by banks, securities firms, and other trading organizations. Since the error analysis is more straightforward for the integration problem of estimating percentage given the loss level than for the rank statistics of estimating loss given a probability level, we study the variance of the probability p that the loss value is greater than a given loss L_q, where L_q is a quantile value at probability level q calculated from standard Monte Carlo simulation in the following numerical study. The formula of the probability can be expressed as

$$P(L_q) = \int 1(L > L_q) f(e^*) de^*, \tag{4.5}$$

where 1(·) is the indicator function and f(·) is the probability density function of the asset return.

4.1.2 Importance Sampling Method

From equations (4.1) and (4.4), almost all the input parameters in the portfolio loss model are generated randomly; hence even for a small number of risk exposures such as 100 loans, it would take a long time for the standard Monte Carlo simulation to estimate the loss distribution. Therefore, we want to make use of some variance reduction methods to speed up the computation. One common method is importance sampling. The idea behind importance sampling is to concentrate simulation on sample paths that contribute most to estimating the expected value; in our case, since the default is a rare event, the importance sampling technique changes the distribution from which the random samples are drawn so that a high number of defaults will occur and result in large losses under the new measure. It has been applied to increase the accuracy and reduce the variance of estimator in the aforementioned default risk portfolio model by Morokoff (2004).

The importance sampling method in Morokoff's model considers a single dimension only and manages to find one which has the largest impact on the portfolio value. For a normal distribution, the usual importance sampling technique used when more samples in the tails are needed is to scale up the variance so that there are more points further out. Specifically, more asset return values in the tails are produced so that more risky assets are likely to default and more losses will accumulate. In

contrast to a simple shift on the covariance matrix, the basic idea here is to orthogonalize the correlation matrix P (equation (4.3)) and scale up the variance corresponding to the largest eigenvalue. The detailed implementation work as follows.

Let Q be the orthogonal matrix whose columns are the orthonormal eigenvectors of correlation matrix P, and Λ be the diagonal matrix of eigenvalues sorted such that $\lambda_1 > \lambda_2 > \ldots > \lambda_N$. Then we have $P = Q\Lambda Q^T$, and the probability density function of asset return vector e, which follows a normal distribution with mean zero and covariance matrix P is

$$f^*(e) = \frac{1}{\sqrt{(2\pi)^N |P|}} \exp(-\frac{1}{2} e^T P^{-1} e).$$

By applying the importance sampling method, we change the correlation matrix $P \to \widetilde{P}$, with $\widetilde{P} = Q\widetilde{\Lambda}Q^T$, where $\widetilde{\lambda}_1 = \theta^2 \lambda_1$ and $\widetilde{\lambda}_j = \lambda_j$, for j > 1. In this way, $\widetilde{P}$ is the covariance matrix which results from scaling up the largest eigenvalue by a factor θ^2. The asset return vector e under this new correlation matrix has the p.d.f. as

$$f(e) = \frac{1}{\sqrt{(2\pi)^N |\widetilde{P}|}} \exp(-\frac{1}{2} e^T \widetilde{P}^{-1} e),$$

and from the properties of orthogonal matrices, we have the weight function of

$$\begin{aligned} w(e) &= \frac{f^*(e)}{f(e)} \\ &= \sqrt{\frac{|\widetilde{P}|}{|P|}} \exp(-\frac{1}{2} e^T [P^{-1} - \widetilde{P}^{-1}] e) \\ &= \theta \exp[-\frac{1}{2}(1 - \frac{1}{\theta^2}) \frac{(q_1^T \cdot e)^2}{\lambda_1}] \end{aligned} \tag{4.6}$$

where q_1 is the first column of Q corresponding the largest eigenvalue λ_1. The asset return generated from importance sampling approach is

$$e = (Q\Sigma Q^T)e^*, \tag{4.7}$$

where Σ is a diagonal matrix with $\Sigma_{11} = \theta$, and $\Sigma_{jj} = 1$ for $j > 1$. In practical implementation, we can generate the asset return e following new correlation matrix and compute weight w(e) according to equation (4.6) after we obtain λ_1 and q_1. The expected loss estimate via importance sampling becomes

$$\begin{aligned} E[L(e^*)] &= \int L(e) f^*(e) de \\ &= \int L(e) \frac{f^*(e)}{f(e)} f(e) de \\ &= \int L(e) w(e) f(e) de \\ &= E[L(e) \bullet w(e)]. \end{aligned}$$

Similarly, the probability that loss exceed some certain quantile level is given by

$$P(L_q) = \int 1(L(e) > L_q) w(e) f(e) de .$$

4.1.3 Importance Sampling Measure Parameter

The fundamental idea of importance sampling is to express an expectation under one probability measure by one under another probability measure through the Radon-Nikodym theorem. The right choice of the new probability measure will result in effective variance reduction. The issue on finding the optimal scale parameter θ that determines the performance of the importance sampling method for the portfolio of credit risky assets has been raised in Morokoff's paper. Morokoff uses trial and error to generate some general guidelines; however, no further study on the optimal

value has been conducted. The efficiency of Monte Carlo simulation depends on the variance of the estimation; hence our goal in this essay is to obtain the optimal choice of θ that minimizes the variance of the loss estimator under new measure with the importance sampling implementation.

The portfolio loss estimator under original measure Q is obtained by generating the asset return based on the original correlation matrix and then taking the sample mean over replications of L(e*), i.e., $\hat{L} = E^{Q}[L(e^*)]$. If there is a new measure P such that measure Q is absolutely continuous with respect to the measure P, then by the Radon-Nikodym theorem, we have

$$\begin{aligned}\hat{L} &= E^{Q}[L] \\ &= \int L dQ \\ &= \int L \frac{dQ}{dP} dP \\ &= E^{P}[L \tfrac{dQ}{dP}],\end{aligned}$$

and the formula $\hat{L} = E^{P}[L(e)\frac{dQ}{dP}]$ gives an alternative unbiased estimator for the loss under measure P. The variance of this new estimator is given by

$$E^{P}[(L\tfrac{dQ}{dP})^2] - (E^{P}[L\tfrac{dQ}{dP}])^2 = E^{P}[(L\tfrac{dQ}{dP})^2] - \hat{L}^2 .$$

Since $\hat{L}^2$ is a constant, the optimization problem can be expressed as follows:

$$\text{Min } V^{P}(\theta) = E^{P}[L^2(e) \bullet (dQ/dP)^2], \tag{4.8}$$

where Q is the original measure and P is the new measure with importance sampling. dQ/dP can be explicitly expressed as w(e) in our model. Notice that the optimization function of importance sampling measure discussed here and the optimal parameter used in the following implementations are all with respect to the expected loss.

Unexpected loss is driven by the same default events and we didn't run the optimization based on quantiles.

4.2 Method Descriptions

4.2.1 IPA_Q Method

We use gradient-based stochastic approximation to solve the variance minimization problem as $\theta^* = \arg\min_{\theta \in \Theta} V(\theta)$. The direct differentiating of the term inside equestion (4.7) requires the derivative of L(e) and dQ/dP. The derivative is complicated, because L(e) is calculated under measure P and is obviously dependent on θ. Su and Fu (2002) propose an optimal importance sampling method, called IPA_Q here, which carries out the simulation to find the best change of measure under original measure Q based on an estimate of the gradient of the variance. It is different from the IPA_P method described by Vazquez-Abad et al. (1998), in that the gradient estimate under the new measure P is used to find the minimized variance.

First, we have

$$
\begin{aligned}
V^P(\theta) &= E^P[L^2(e)(\frac{dQ}{dP})^2] \\
&= \int_\Omega L^2(e)\frac{(dQ)^2}{(dP)^2}dP \\
&= \int_\Omega L^2(e)\frac{(dQ)^2}{dP} \\
&= \int_\Omega L^2(e)\frac{dQ}{dP}dQ \\
&= E^Q[L^2(e)\frac{dQ}{dP}] \\
&= V^Q(\theta).
\end{aligned}
\tag{4.9}
$$

Therefore, the importance sampling problem is transformed into a minimization problem under original measure, which eliminates the dependence between the loss function and the parameters in the optimization. Recall that dQ/dP is the weight function, the minimization problem in Equation (4.8) then becomes

$$\text{Min } V^{Q}(\theta) = E^{Q}[L^2(e^*) \bullet w(e^*, \theta)]. \tag{4.10}$$

Note that the loss function L(e*) does not depend on θ under measure Q.

We next derive the IPA_Q gradient estimator as follows. According to the portfolio model, we have

$$w(e^*, \theta) = \theta \exp[-0.5(1 - \frac{1}{\theta^2}) \frac{(q_1^T \cdot e^*)^2}{\lambda_1}].$$

Then,

$$\begin{aligned} \frac{\partial w(e^*, \theta)}{\partial \theta} &= (1 - \frac{1}{\theta^2} \frac{(q_1^T \cdot e^*)^2}{\lambda_1}) \exp[-0.5(1 - \frac{1}{\theta^2}) \frac{(q_1^T \cdot e^*)^2}{\lambda_1}] \\ &= (1 - \frac{M}{\theta^2}) \exp[-0.5(1 - \frac{1}{\theta^2})M]. \end{aligned}$$

Defining $M := \frac{(q_1^T \cdot e^*)^2}{\lambda_1} > 0$ for notional convenience, the gradient estimator becomes

$$\begin{aligned} g(\theta) &= \frac{\partial V^{Q}(\theta)}{\partial \theta} \\ &= E^{Q}[L^2 \bullet \frac{\partial w(\theta)}{\partial \theta}] \\ &= E^{Q}[L^2 \bullet (1 - \frac{M}{\theta^2}) \exp[-0.5(1 - \frac{1}{\theta^2})M]]. \end{aligned} \tag{4.11}$$

The derivation shows that IPA_Q method provides a much simpler gradient estimator, since no differentiability of the loss function with respect to the optimization parameter θ is needed.

Next we show the objective function is a convex function, so that the IPA estimator is unbiased and the simulation converges to the optimum. The second derivative of objective function to parameter is

$$\frac{\partial^2 V^Q(\theta)}{\partial \theta^2} = E^Q[L^2 \bullet \frac{\partial^2 w(\theta)}{\partial \theta^2}]$$
$$= E^Q[L^2 \bullet (M^2\theta^{-5} + M\theta^{-3}) \exp[-0.5(1-\frac{1}{\theta^2})M]] > 0.$$

The last inequality is valid since $L^2 > 0$, $\exp(\bullet) > 0$, and $M^2\sigma^{-5} + M\sigma^{-3} > 0$. This positive second derivative guarantees an optimum exists for the objective function.

After the IPA estimator is obtained, we can find the optimal parameter θ* for importance sampling procedure via an iterative scheme:

$\theta_{k+1} = \theta_k - a_k g_k$,

where θ_k is the parameter at k^{th} iteration, g_k is an estimator of the derivative with respect to θ of the variance of the estimator V(θ), and $\{a_k\}$ is a positive sequence of numbers converging to 0. The selection of the step size is critical and we choose the step size according to varying initial conditions by trial and error. The simulation stops when $|a_k g_k| < 10^{-5}$.

The algorithm for applying importance sampling via IPA_Q is as follows:

Algorithm IPA_Q

Stage I: Find θ*.

1. Initialize: set $\theta = \theta_0$, $k = 1$.
2. Repeat until a specified stopping rule is satisfied:
   a. Generate sample paths under original measure Q.
   b. Calculate IPA_Q and g_k based on equation (4.11).

c. Update $\theta_{k+1} = \theta_k - a_k g_k$.

d. Set k ← k+1.

3. Set $\theta^* = \theta_{k+1}$.

Stage II: Simulate the loss model using importance sampling at θ = θ*.

4.2.2 Cross-entropy Method

Since the selection of importance sampling measure change parameter is a minimization problem as showed in equation (4.8), we can apply the CE method to obtain the best solution. The CE method adaptively finds the optimal value. We assume the importance sampling measure parameter to be estimated in the optimization problem is generated from a normal distribution model at each iteration. We will compare the performance of IPA_Q and CE algorithms in the following numerical examples.

Algorithm CE

Stage I: Find θ*.

1. Initialize: Specify quantile parameter ρ and sample size N. Initialize parameters of the probabilistic model (normal distribution) μ_0 and σ_0. Set k=0.
2. Repeat until a specified stopping rule is satisfied:
    a. Generate N i.i.d. samples $\theta_{1,k}, \ldots, \theta_{N,k}$ from the $N(\hat{\mu}_k, \hat{\sigma}_k)$ distribution.
    b. Select the ρN best performing (elite) samples, and let I be the indices of the ρN best performing samples.
    c. Update the parameters as:

$$\mu_{k+1} = \frac{1}{\rho N}\sum_{i\in I}\theta_{i,k}$$, and

$$\sigma_{k+1} = \frac{1}{\rho N}\sum_{i\in I}(\theta_{i,k} - \mu_{k+1})^2$$

d. Smooth by using equation (3.6), (3.7) and (3.8).

e. Set k ← k+1.

3. Set $\theta^* = \mu_{k+1}$.

Stage II: Simulate the loss model using importance sampling at θ = θ*.

4.3 Numerical Analysis

In this session, we apply both IPA_Q and CE methods in a test case to illustrate the effectiveness of the methods in variance reduction of importance sampling in the Monte Carlo simulation. Our goal is to find the minimum $V^Q(\theta)$ as shown in the equation (4.10). In our numerical example, we consider a fixed horizon of one year and have the following input parameters:

The number of risky assets in the portfolio is N = 100.

The number of factors in equation (4.2) is M = 50.

The factor z_i is generated from the standard normal distribution.

The factor loadings $B = [\beta_1, \beta_2, \ldots, \beta_M]$ are generated randomly. The weight corresponding to the first factor is uniformly distributed on [0.21, 0.31], the weights on the second to fifth factors are uniformly distributed on [0.11, 0.21], and two additional factors are chosen randomly from the remaining 45 factors, and their

weights are uniformly distributed on [0, 0.1]. All other factors loadings are zero. Finally, the factor loadings are normalized so that $\beta_i^T \beta_i = 1$.

R-square r_i^2 is uniformly distributed on [0.1, 0.4].

The default probability has the formula of $p_i = 0.01(1/\sqrt{r_i^2} - 1)$, and the default threshold $\alpha_i = \Phi^{-1}(p_i)$, where Φ is the cumulative density function of the standard normal distribution.

ε_i is generated from i.i.d. standard normal distribution.

LGD_i is generated from Beta distribution with mean 0.5 and standard deviation 0.25.

We run 1,000 replications to find the optimal importance sampling parameter based on the expected loss, and the results are shown in Figures 4.1 and 4.2. For the IPA_Q method, we conduct the experiments with initial θ_0 = 1, 2, 3, 4, 5, 6, 7, 8, 9, 10. Figure 4.2 gives the results of the CE algorithm with different values of the initial condition. We assume θ is generated from a normal distribution model at each iteration, and two sets of initial conditions are tested: initial mean of the normal distribution is 5.0 with initial standard deviation of 3.0; and initial mean of 10.0 with initial standard deviation of 5.0. The above two lines in Figure 4.2 shows the evolution of θ under varying initial conditions, and the bottom two lines illustrate the evolution of standard deviation as simulation progresses. The standard deviation in the preset normal distribution model reaches zero. Using the CE method, convergence is achieved within the first five or six iterations of the stochastic simulation. Both methods converge to the optimum quickly regardless of the initial choice and give an

optimal value θ around 1.5. The running time was 45 seconds per iteration for IPA_Q method and 15 minutes per iteration for CE method on a 1.5GHz Pentium PC; therefore the total computational time to achieve the convergence for IPA_Q method is less than that of CE method.

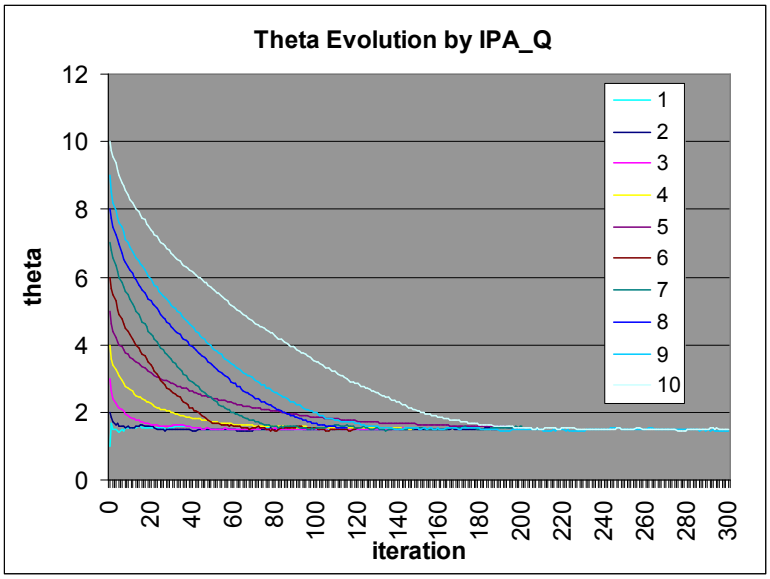

Figure 4.1 Optimal theta evolution by IPA_Q for EL

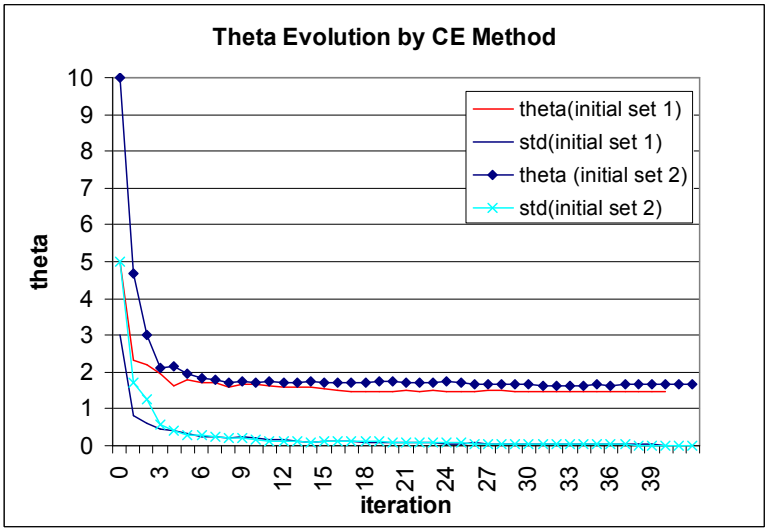

Figure 4.2 Optimal theta evolution by CE method for EL

In our numerical case, we have the exact formula of expected loss as $EL = \overline{LGD} * \overline{p}$. Considering LGD following a Beta distribution with standard deviation of 0..25 and p is a function of R-square which is uniformly distributed between 0.1 and 0.4, we can calculate the analytical value of the expected loss as 0.5541%. We apply the optimal IS scalar parameter obtained by both methods, 1.5, to the simulation based on 100,000 replications. Table 4.1 compares the simulation error of EL and UL for the standard Monte Carlo simulation, the Monte Carlo simulation with importance sampling method by selecting $\theta = 3$ obtained from Morokoff's trial and error process, and the IS method by choosing $\theta = 1.5$ from our optimization results. In the table, σ_s represents the standard deviation of the raw Monte Carlo simulation, $\sigma_{IS\text{-}M}$ is the standard deviation of the IS method with Morokoff's scale parameters, and σ_{IS} the standard deviation of the IS method with scalar from IPA_Q and CE methods. In the table, we also show the ratio of the estimate of naïve variance and the estimate of importance sampling variance, which is sometimes called variance reduction factor. We found our importance sampling approach leads to a reduction of expected loss standard deviation from 1.06% to 0.75%. Substantial improvement in accuracy is found in estimating unexpected loss with the standard deviation decreasing from 0.074% to 0.017%. Because the importance sampling method requires only a small amount of additional computation to evaluate asset return and weight under new measure, the variance reduction factor is a measure of the degree of computational savings achieved by importance sampling. For our experiment, we found that the expected loss could run 2 times faster, while for the unexpected loss the speed up is 18. According to the variance ratio in the table, the

same accuracy in estimating EL obtained by a standard Monte Carlo simulation at 100,000 runs could be acquired with the Morokoff importance sampling method using approximately 8550 runs, or with our importance sampling method using around 50,000 runs. The accuracy enhancement becomes even more pronounced in estimating the unexpected loss. 100,000 runs of standard Monte Carlo simulation is reduced to about 24,500 runs in Morokoff IS method, and 5500 runs in our method. The variance estimator of the IS method with scalar 1.5 is about 42% and 78% lower than the calculation with scalar 3, in estimating the expected loss and unexpected loss, respectively. The results indicate that our optimization algorithm gives a better importance sampling measure parameter than the one from Morokoff's trial and error method. For instance, in order to achieve the same level of precision as IPA_Q in estimating expected loss, Morokoff's method would require approximately 1.7 times as many simulations, while for unexpected loss it is about 4.5 times.

Other than the expected loss and unexpected loss functions, we also examine the accuracy improvement of the simulation in the tail quantile levels. According to equation (4.5), the standard deviation σ_s without importance sampling is

$$\begin{aligned}\sigma_s^2 &= \int (1(L(e) > L_q) - q)^2 f^*(e) de \\ &= \sqrt{q(1-q)}.\end{aligned}$$

A similar formula is obtained for the importance sampling simulation as

$$\sigma_{IS}^2 = \int (1(L(e) > L_q) \cdot w(e) - q)^2 f(e) de .$$

Table 4.1 EL and UL for portfolio of 100 exposures

		Expected Loss (%)	Unexpected Loss (%)
True Value	Mean	0.5541	NA
Standard MCS	Mean	0.5534	1.0616
	σ_s	1.0616	0.0739
IS – Morokoff (θ=3.0)	Mean	0.5561	0.9817
	$\sigma_{IS\text{-}M}$	0.9817	0.0367
IS (θ=1.5)	Mean	0.5533	0.7503
	σ_{IS}	0.7503	0.0174
$\sigma_s^2/\sigma_{IS\text{-}M}^2$		1.17	4.05
σ_s^2/σ_{IS}^2		2.00	18.04

In the same way as we treat the expected loss, we use IPA_Q method to find the optimal importance sampling measure for the probability that loss exceed some certain quantile level. The minimization problem becomes

Min $V^Q(\theta) = E^Q[1(L > L_q)^2 \bullet w(e^*, \theta)]$,

and the gradient estimator is

$$g(\theta) = \frac{\partial V^Q(\theta)}{\partial \theta}$$
$$= E^Q[1(L > L_q)^2 \bullet (1 - \frac{M}{\theta^2}) \exp[-0.5(1 - \frac{1}{\theta^2})M]].$$

We also see that the objective function $V^Q(\theta)$ is a convex function by calculate the second derivative with respect to θ.

Because we use IPA under the original measure, we only need the loss at some certain quantile level obtained from a standard Monte Carlo simulation. From our calculation, L_q = 5.06%, 10.01%, and 14.20% corresponding to q = 1%, 0.1%, and 0.01%, respectively, as shown in table 4.2. Figure 4.3 shows the evolution of finding optimal θ for q = 1% with initial condition of θ_0 = 1 and 5. Both optimization converge to θ* = 2.3. By applying the same approach, we find the optimal IS measure is θ* = 2.5 for q = 0.1%, and θ* = 5.0 for q = 0.01%. In the following experiments, we will use those optimal values in our importance sampling simulation to estimate loss at different quantile levels.

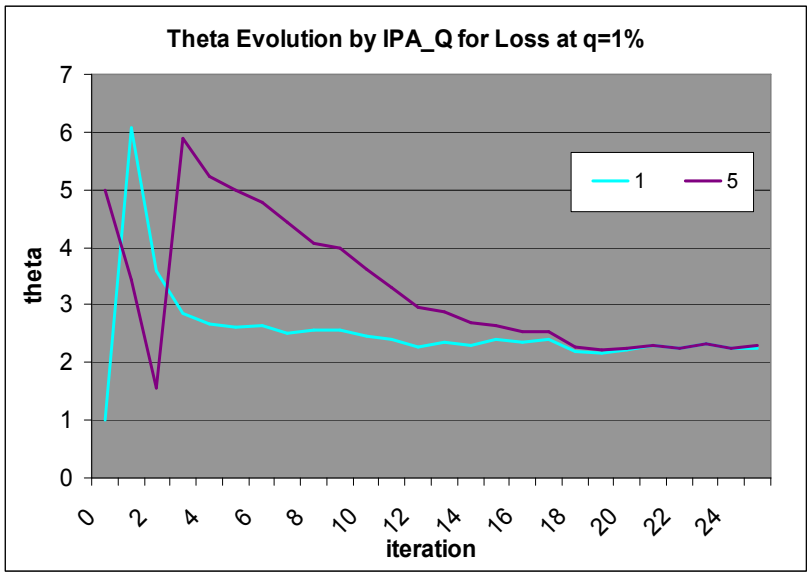

Figure 4.3 Optimal theta evolution by IPA_Q method for loss at q= 1%

Table 4.2 shows the quantile results of the credit risk portfolio at the level of 1%, 0.1%, and 0.01%, based on 10,000 runs. From the standard Monte Carlo

simulation, the quantile of the loss distribution is 0.0506 at quantile level 1%, 0.1001 at 0.1%, and 0.1420 at level 0.01%. The large loss value indicates the loss distribution is a fat tailed distribution. The results show that a significant improvement in accuracy is achieved by applying importance sampling method and more improvement is obtained by using the optimal IS measure parameter. At the 1% level, the standard Monte Carlo method gives σ_s = 9.95%, the Morokoff importance sampling gives $\sigma_{IS\text{-}M}$ = 4.97%, and our method has σ_{IS} = 4.43%. In terms of the computational speed up, there is a speed up factor of 5.04 at level 1%, 27.24 at 0.1% level, and 82.64 at 0.01% level by applying the optimal measure scalar. As shown in the table, the results of $\sigma_{IS\text{-}M}$ and σ_{IS} at different quantile levels indicate that the variance is reduced effectively by using the optimal importance sampling measure parameter acquired from IPA_Q and CE methods.

Table 4.2 Quantiles for Portfolio of 100 exposures

Quantile level q / Loss		1 %	0.1%	0.01%
Loss L_q		5.06%	10.01%	14.20%
Standard MCS	σ_s	9.95%	3.16%	1.00%
Morokoff IS (θ=3.0)	$\sigma_{IS\text{-}M}$	4.97%	0.61%	0.13%
IS	σ_{IS}	4.43%	0.60%	0.11
$\sigma_s^2/\sigma_{IS\text{-}M}^2$		4.01	26.84	59.17
σ_s^2/σ_{IS}^2		5.04	27.74	82.64

4.4 Antithetic Variate Method

Antithetic variate (AV) approach is a common variance reduction technique for increasing the precision of the estimates, and it usually accelerates the convergence. This method is easy to apply, because it concentrates on the process used for generating the random deviates. The fundamental idea behind is to bring in negative correlation between two estimates.

Suppose that a sequence of random variates are generated using the random number sequence $u_1, u_2, \ldots, u_n$, where the $\{u_i\}$ are uniformly distributed in the interval (0,1). Then the sequence $(1- u_1)$, $(1- u_2)$,…, $(1- u_n)$ follow the same uniform distribution and are used to generate a second set of variates. Assume the estimate of the value function using the first set is g(u), and the estimate using the second set is g(1-u), then we have an unbiased estimate of

$$\frac{1}{2}[g(u)+g(1-u)],$$

with variance

$$\frac{1}{4}[\text{var}(g(u))+\text{var}(g(1-u))]+\frac{1}{2}[\text{cov}(g(u),g(1-u))].$$

If the covariance between g(u) and g(1-u) is negative, this will yield a smaller estimate of the variance than an independent estimate.

For those standard normally distributed random variables used in our model, we first generate a set of random normal deviates for the initial estimate, and a second estimate is then obtained by using the same set of random normal deviates with their signs reversed because if x is a standard normal random variable, so is –x.

Table 4.3 shows the variance reduction results by applying antithetic variate method based on 10,000 replications. We estimate the expected loss, unexpected loss and quantiles. Compared with standard Monte Carlo simulation, the variance is reduced approximately by 10% by applying antithetic sampling only. The speedup factor of simulation is in the range between 1.00 and 1.66. We try to combine antithetic variate method with importance sampling to reduce the variance further, however, there is no obvious improvement over the results by using importance sampling only, as illustrated in table 4.1 and 4.2. The results indicate that the effect of importance sampling in reducing variance dominates that of antithetic sampling.

From the outcome we find the speed improvement of using antithetic variate approach is minor in our model. The reason behind is that although normal variates have perfect negative correlations, this does not hold for the corresponding transformed return function in the portfolio loss model. Therefore, while the covariance term is negative, its magnitude is not large enough to result in a significant reduction in the variance of the revised estimate.

Table 4.3 Variance Reduction by Using Antithetic Variate Method

	EL	UL	P(0.1)	P(0.01)	P(0.001)
σ_s	1.07%	0.078%	9.95%	3.16%	1.00%
σ_{AV}	1.03%	0.071%	9.29%	2.45%	1.00%
$\sigma_{AV\text{-}IS}$	0.76%	0.018%	4.45%	0.62%	0.12
σ_s^2/σ_{AV}^2	1.08	1.21	1.15	1.66	1.00
$\sigma_s^2/\sigma_{AV\text{-}IS}^2$	1.98.	18.8	5.00	25.98	69.44

4.5 Conclusions

This essay describes two stochastic optimization methods for choosing an optimal importance sampling measure change factor, thereby improving the accuracy of a Monte Carlo simulation used to estimate the loss distribution on a portfolio of credit risky assets. Scaling up the scalar parameter in the asset correlation model increases the correlations, thus inducing a larger number of correlated defaults and generating samples further out in the loss tail. The implementation of importance sampling with optimal scalar yields more precise estimates of the descriptive statistics value than naïve simulation. It leads to a significant variance reduction in estimating expected loss, loss volatility and quantile. For example, in the case of estimating the volatility, it reduces the number of simulation runs by a factor of 18 compared with a standard Monte Carlo simulation and 4.5 compared with Morokoff's choice of importance sampling measure. More substantial variance reduction is expected if more risky assets are involved.

When using the gradient-based method to estimate the optimal importance sampling measure, we cast the minimization problem under the original probability measure, which removes the dependence between the loss function and the scalar parameter in the optimization. As a result, we do not require differentiability of the loss function with respect to the underlying variable and our method is applicable in much more general settings. It results in a simpler IPA gradient estimator than the original IPA estimator. We further prove that the objective function in our minimization problem is a convex function. We also apply antithetic variate method to further reduce variance and the effect of antithetic sampling is limited.

Future studies include applying copula function in the portfolio analysis to account for the credit risk dependence structure. In addition, we can relax the assumption of the multivariate normal distribution for the asset return because our approach could apply to a multivariate distribution for which the correlation matrix is used in the sampling process.

Bibliography

E. Aarts, and J. Korst, *Simulated Annealing and Boltzmann Machines*, John Wiley & Sons, 1989.

J. Barraquand, and D. Martineau, "Numerical Valuation of High Dimensional Multivariate American Securities," *Journal of Financial and Quantitative Analysis*, 30, 383-405, 1995.

F. Black, and J. Cox, "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions," *Journal of Finance*, 31, 351-367, 1976.

F. Black, and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654, 1973.

M. Broadie, and P. Glasserman, "Pricing American-Style Securities Using Simulation," *Journal of Economic Dynamics and Control*, 21, 1323-1352, 1997.

M. Broadie, and P. Glasserman, "A Stochastic Mesh Method for Pricing High-Dimensional American Options," *Journal of Computational Finance*, 7, 35-72, 2004.

J. F. Carriere, "Valuation of the Early-Exercise Price for Derivative Securities using Simulations and Splines," *Insurance: Mathematics and Economics,* 19, 19-30, 1996.

X. Chen, and D. Simchi-Levi, "Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case," *Operations Research***,** 52, 887-896, 2004.

A. Clark, and H. Scarf, "Optimal Policies for a Multi-Echelon Inventory Problem," *Management Science*, 6, 475-490, 1960.

G. Dueck, and T. Scheur, "Threshold Accepting: A General Purpose Optimization Algorithm Appearing Superior to Simulated Annealing," *Journal of Computational Physics*, 90, 161-175, 1990.

D. Duffie, and M. Huang, "Swap Rates and Credit Quality," *Journal of Finance*, 51, 921-949, 1996.

D. Duffie, and K. Singleton, "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies*, 12, 687-720, 1997.

A. Federgruen, and A. Heching, "Combined Pricing and Inventory Control Under Uncertainty," *Operations Research*, 47, 454-475, 1999.

Y. Feng, and B. Xiao, "A Continuous-Time Yield Management Model with Multiple Prices and Reversible Price Changes," *Management Science*, 46, 644-657, 2000.

M. Fu, "Sample Path Derivatives for (s, S) Inventory Systems," *Operations Research*, 42, 351-364, 1994a.

M. Fu, "Optimization Using Simulation: A Review," *Annals of Operations Research*, 53, 199-248, 1994b.

M. Fu, "Optimization for Simulation: Theory vs. Practice," *INFORMS Journal on Computing*, Vol.14, No.3, 192-215, 2002.

M. Fu, "Gradient Estimation," Chapter 19 in *Handbooks in Operations Research and Management Science*: *Simulation*, S.G. Henderson and B.L. Nelson, eds., Elsevier, 575-616, 2006.

M. Fu, and K. Healy, "Techniques for Simulation Optimization: An Experimental Study on an (s, S) Inventory System," *IIE Transactions*, Vol.29, No.3, 191-199, 1997.

M. Fu, and J. Hu, "Sensitivity Analysis for Monte Carlo Simulation of Option Pricing," *Probability in the Engineering and Information Sciences*, 9, 417-446, 1995.

M. Fu, and J. Hu, *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*, Kluwer Academic Publishers, 1997.

M. Fu, S. Laprise, D. Madan, Y. Su, and R. Wu, "Pricing American Options: A Comparison of Monte Carlo Simulation Approaches," *Journal of Computational Finance*, 4, 39-88, 2001.

C. Gallego, and G. Van Ryzin, "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, 40, 999-1020, 1994.

P. Glasserman, *Gradient Estimation Via Perturbation Analysis*, Kluwer Academic Publishers, 1991.

P. Glasserman, *Monte Carlo Methods in Financial Engineering*, Springer Verlag, New York, 2003.

P. Glasserman, and J. Li, "Importance Sampling for Portfolio Credit Risk," *Management Science*, 51, 1643-1656, 2005.

F. Glover, "Tabu Search: A Tutorial," *Interfaces*, 20, 74-94, 1990.

D. Grant, G. Vora, and D. Weeks, "Simulation and the Early-Exercise Option Problem," *Journal of Financial Engineering*, 5, 211-227, 1996.

D. Grant, G. Vora, and D. Weeks, "Path-Dependent Options: Extending the Monte Carlo Simulation Approach," *Management Science*, 43, 1589-1602, 1997.

Y. C. Ho, and X. R. Cao, *Discrete Event Dynamic Systems and Perturbation Analysis*, Kluwer Academic, 1991.

R. H. Hollier, K. L. Mak, and K. Yiu, "Optimal Inventory Control of Lumpy Demand Items Using (*s*, *S*) Policies with a Maximum Issue Quantity Restriction and Opportunistic Replenishments," *International Journal of Production Research*, 43, 4929 – 4944, 2005.

J. Hu, M. Fu, and S. Marcus, "A Model Reference Adaptive Search Algorithm for Global Optimization," *Operations Research*, 2007.

J. Hull, and A. White, "Efficient Procedures for Valuing European and American Path-Dependent Options," *J. Derivatives*, 1, 21-31, 1993.

D. Iglehart, "Optimality of (s, S) Policies in the Infinite Horizon Dynamic Inventory Problem," *Management Science*, 9, 259-267, 1963.

S. Jacobson, and L. Schruben, "A Review of Techniques for Simulation Optimization," *Operations Research Letters*, 8, 1-9, 1989.

R. Jarrow, D. Lando, and S. Turnbull, "A Markov Model for the Term Structure of Credit Spreads," *Review of Financial Studies*, 10, 481-523, 1997.

R. Jarrow, and S. Turnbull, "Pricing Derivatives on Financial Securities Subject to Credit Risk," *Journal of Finance*, 50, 53-86, 1995.

M. Kalkbrener, H. Lotter, and L. Overbeck, "Sensible and Efficient Capital Allocation for Credit Portfolios," *Risk*, 2004.

S. Kealhofer, and J. Bohn, "Portfolio Management of Default Risk," available online via *www.moodyskmv.com*, 2001.

C. Keber, "Option Pricing with the Genetic Programming Approach," *Journal of Computational Intelligence in Finance*, 7, 26-36, 1999.

S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by Simulated Annealing," *Science*, 220, 671-680, 1983.

D. Kroese, R. Rubinstein, and S. Porotsky, "The Cross-Entropy Method for Continuous Multiextremal Optimization," *Operations Research*, 2004.

H. Kushner, and G. Yin, *Stochastic Approximation Algorithms and Applications,* Springer-Verlag, New York, 1997.

S. Laprise, M. Fu, S. Marcus, A. Lim, and H. Zhang, "Pricing American-style Derivatives with European Call Options," *Management Science*, 52, 95-110, 2006.

P. Larranaga, R. Etxeberria, J. Lozano, B. Sierra, I. Inza, and J. Pena, "A Review of the Cooperation Between Evolutionary Computation and Probabilistic Graphical Models," *Proceedings of the Second Symposium on Artificial Intelligence. Adaptive Systems. CIMAF 99. Special Session on Distributions and Evolutionary Computation,* 314-324, 1999.

P. L'Ecuyer, R. Simard, E. Chen, and W. Kelton, "An Objected-Oriented Random-Number Package with Many Long Streams and Substreams," *Operations Research,* 50, 1073—1075, 2002.

F. Longstaff, and E. Schwartz, "Valuing American Options by Simulation: A Simple Least-Squares Approach," *The Review of Financial Studies*, 14, 113-148, 2001.

D. Madan, and H. Unal, "Pricing the Risks of Default," *Review of Derivatives Research*, 2, 121-160, 1998.

H. Markowitz, "Portfolio Selection," *Journal of Finance*, 7, 77-91, 1952.

R. Merton, "On the Pricing of Corporate Debt: the Risk Structure of Interest Rates," *Journal of Finance*, 29, 449-470, 1974.

R. Merton, “Option Pricing when Underlying Stock Returns are Discontinuous,” *Journal of Financial Economics*, 3, 125-144, 1976.

W. Morokoff, “An Importance Sampling Method for Portfolios of Credit Risky Assets,” *Proceedings of Winter Simulation Conference*, 2004.

H. Mühlenbein, and G. Paaβ, “From Recombination of Genes to the Estimation of Distributions: I. Binary Parameters,” *In Lecture Notes in Computer Science 1411: Parallel Problem Solving from Nature – PPSN IV*, 178-187, 1996.

J. Nelder, and R. Mead, “A Simplex Method for Function Minimization,” *Computer Journal*, 7, 308-313, 1965.

T. Paul, and H. Iba, “Linear and Combinatorial Optimizations by Estimation of Distribution Algorithms,” *9th MPS Symposium on Evolutionary Computation*, *IPSJ*, Japan, 2002.

H. Robbins, and S. Monro, “A Stochastic Approximation Method,” *Annals of Mathematical Statistics*, 22, 400-407, 1951.

R. Rubinstein and D. Kroese, *The Cross-entropy Method: a Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning*, Springer-Verlag, New York, 2004.

M. Safizadeh, “Optimization in Simulation: Current Issues and the Future Outlook,” *Naval Research Logistics,* 37, 807-825, 1990.

H. Scarf, “The Optimality of (s, S) Policies in they Dynamic Inventory Problem,” *Mathematical Methods in the Social Sciences*, Stanford University Press, 1960.

H. Scarf, “Inventory Theory,” *Operations Research*, 50, 186-191, 2002.

J. Spall, “Multivariate Stochastic Approximation Using a Simultaneous Perturbation Gradient Approximation,” *IEEE Transactions on Automatic Control*, 37, 332-341, 1992.

J. Spall, “An Overview of the Simultaneous Perturbation Method for Efficient Optimization,” *Johns Hopkins APL Technical Digest*, 19, 482–492, 1998.

J. Spall, “Stochastic Optimization, Stochastic Approximation and Simulated Annealing,” in *Encyclopedia of Electrical and Electronics Engineering*, J. G. Webster, ed., Wiley, New York, 20, 529-542, 1999.

J. Spall, *Introduction to Stochastic Search and Optimization*, Wiley, Hoboken, NJ, 2003.

M. Srinivas, and L. Patnaik, “Genetic Algorithms: A Survey,” *IEEE Computer*, 27, 17-26, 1994.

Y. Su, and M. Fu, “Optimal Importance Sampling in Securities Pricing,” *Journal of Computational Finance*, 5, 2002.

L. Thomas, “Price and Production Decision with Random Demand,” *Operations Research*, 22, 513-518, 1974.

J. Tilley, “Valuating American Options in a Path Simulation Model,” *Transactions of the Society of Actuaries*, 45, 83-104, 1993.

J. Tsitsiklis, and B. Van Roy, “Regression Methods for Pricing Complex American-style Options,” *IEEE Transactions on Neural Networks*, 14, 694-703, 2001.

F. Vazquez-Abad, and D. Dufresne, “Accelerated Simulation for Pricing Asian Options,” *Proceedings of Winter Simulation Conference*, Washington, D.C., 1493-1500, 1998.

A. Veinott, and H. Wagner, "Computing Optimal (s, S) Inventory Policies," *Management Science*, 11, 525-552, 1963.

A. Zhigljavsky, *Theory of Global Random Search*, Kluwer Academic, Boston, 1991.