Risk Aversion in Inventory Management*

Xin Chen†, Melvyn Sim‡, David Simchi-Levi§ and Peng Sun¶

February 22, 2003

Abstract

Traditional inventory models focus on risk neutral decision makers, i.e., characterizing replenishment strategies that maximize expected total profit, or equivalently, minimize expected total cost over a planning horizon. In this paper, we propose a general framework for incorporating risk aversion in multi-period inventory models as well as multi-period models that coordinate inventory and pricing strategies. In each case, we characterize the optimal policy for various measures of risk that have been commonly used in the finance literature. In particular, we show that the structure of the optimal policy for a decision maker with exponential utility function is almost identical to the structure of the optimal risk neutral inventory (and pricing) policies. Computational results demonstrate the importance of this approach not only to risk averse decision makers, but also to risk neutral decision makers with limited information on the demand distribution.

1 Introduction

Traditional inventory models focus on characterizing replenishment policies so as to maximize the expected total profit, or equivalently, to minimize the expected total cost over a planning horizon. Of course, this focus on optimizing expected profit or cost is appropriate for risk neutral decision makers, i.e., inventory mangers that are insensitive to profit variations.

Evidently, not all inventory managers are risk neutral; many planners are willing to tradeoff lower expected profit for downside protection against possible losses. Indeed, some experimental evidence suggests that for some products, the so-called *high-profit* products, the decision makers are risk averse; see Schweitzer and Cachon [27] for more details. Unfortunately, traditional inventory control models fail short of meeting the needs of risk averse planners. For instance, traditional inventory models do not suggest mechanisms to reduce the chance of unfavorable profit levels. Thus, it is important to incorporate the notions of risk aversion in a broad class of inventory models.

The literature on risk averse inventory models is quite limited and mainly focuses on single period problems. Lau [17] analyzes the classical newsvendor model under two different objective functions.

*Research supported in part by the Center of eBusiness at MIT, ONR Contracts N00014-95-1-0232 and N00014-01-1-0146, and by NSF Contracts DMI-9732795, DMI-0085683 and DMI-0245352.

†MIT Operations Research Center, 77 Mass. Ave. E40-110, Cambridge, MA 02139, USA. Email: xchen@mit.edu

‡NUS Business School, National University of Singapore. Email: dscsimm@nus.edu.sg

§MIT Department of Civil and Environmental Engineering, and The Engineering System Division, 77 Mass. Ave. Rm. 1-171, Cambridge, MA 02139, USA. Email: dslevi@mit.edu

¶Fuqua School of Business, Duke University Box 90120, Durham, NC 27708, USA. Email: psun@duke.edu

In the first objective, the focus is on maximizing the decision maker's expected utility of total profit. The second objective function is the maximization of the probability of achieving a certain level of profit.

Eeckhoudt, Gollier and Schlesinger [13] focus on the effects of risk and risk aversion in the newsvendor model when risk is measured by expected utility functions. In particular, they determine comparative-static effects of changes in the various price and cost parameters in the risk aversion setting.

Chen and Federgruen [6] analyze the mean-variance tradeoffs in newsvendor models as well as some standard infinite horizon inventory models. Specifically, in the infinite horizon models, Chen and Federgruen focus on the mean-variance tradeoff of customer waiting time as well as the mean-variance tradeoffs of inventory levels. Martínez-de-Albéniz and Simchi-Levi [21] study the mean-variance tradeoffs faced by a manufacturer signing a portfolio of option contracts with its suppliers and having access to a spot market.

The paper by Bouakiz and Sobel [5] is closely related to ours. In this paper, the authors focus on a decision maker with exponential utility function. The objective is to characterize the inventory replenishment strategy so as to minimize the expected utility of the present value of costs over a finite planning horizon or an infinite horizon. Assuming linear ordering cost, they prove that a base stock policy is optimal.

So far all the papers referenced above assume that demand is exogenous. A rare exception is Agrawal and Seshadri [1] who consider a risk averse retailer which has to decide on its ordering quantity and selling price for a single period. They demonstrate that different assumptions on the demand-price function may lead to different properties of the selling price.

In this paper, we propose a general framework for incorporating risk aversion in multi-period inventory (and pricing) models. We consider two different measures of risk that have been used in the finance literature:

- Increasing concave utility function and its special case, the exponential utility function;
- Conditional-Value-at-Risk, or CVaR.

For each risk measure, we consider two closely related problems. In the first one, demand is exogenous, i.e., price is not a decision variable, while in the second one demand depends on price and price is a decision variable. In both cases, we distinguish between models with fixed ordering costs and models with no fixed ordering cost. For each model, the objective is to find an inventory policy (and a pricing strategy, when price is a decision variable) so as to maximize either the expected utility or the CVaR of wealth at the end of the planning horizon. Observe that if the utility function is linear and increasing, the decision maker is risk neutral and these problems are reduced to the classical finite horizon stochastic inventory problem and the finite horizon inventory and pricing problem.

We summarize known and new results in Table 1.

The row "Risk Neutral Model" presents a summary of known results. For example, when price is not a decision variable, and there exists a fixed ordering cost, $k > 0$, Scarf [26] proved that an (s, S) inventory policy is optimal. In such a policy, the inventory strategy at period t is characterized by two parameters (s_t, S_t). When the inventory level x_t at the beginning of period t is less than s_t, an order of size $S_t - x_t$ is made. Otherwise, no order is placed. A special case of this policy is the base stock policy, in which $s_t = S_t$. This policy is optimal when $k = 0$.

	Price Not a Decision		*Price is a Decision*	
	$k = 0$	$k > 0$	$k = 0$	$k > 0$
Risk Neutral Model	base stock	(s, S)	base stock list price	(s, S, A, p)
Increasing & Concave Utility	wealth dependent base stock	?	wealth dependent base stock	?
Exponential Utility	base stock	(s, S)	base stock	(s, S, A, p)
Conditional Value-at-Risk	wealth dependent base stock	?	wealth dependent base stock	?
Myopic CVaR	base stock	(s, S)	base stock	(s, S, A, p)

Table 1: Summary of Results

If price is a decision variable and there exists a fixed ordering cost, the optimal policy of the risk neutral model is an (s, S, A, p) policy, see Chen and Simchi-Levi [8]. In such a policy, the inventory strategy at period t is characterized by two parameters (s_t, S_t) and a set $A_t \in [s_t, (s_t+S_t)/2]$, possibly empty depending on the problem instance. When the inventory level x_t at the beginning of period t is less than s_t or $x_t \in A_t$, an order of size $S_t - x_t$ is made. Otherwise, no order is placed. Price depends on the initial inventory level at the beginning of the period. A special case of this model is when $k = 0$, for which a base stock list price policy is optimal. In this policy, inventory is managed based on a base stock policy and price is a non-increasing function of inventory at the beginning of each period, see Federgruen and Heching [14] and Chen and Simchi-Levi [8].

The table suggests that when risk is measured using an expected exponential utility function of the final wealth, the structures of optimal policies are almost the same as the one under the risk neutral case. For example, when price is not a decision variable and $k > 0$, the optimal replenishment strategy follows the traditional inventory policy, namely an (s, S) policy. A corollary of this result is that a base stock policy is optimal when $k = 0$. This is exactly the optimal policy characterized by Bouakiz and Sobel [5].

When risk is measured using a general increasing concave utility function or CVaR and there is no fixed ordering cost, the optimal policy is such that inventory is managed using a base stock policy independent of whether price is a decision variable. Unfortunately, this base stock policy depends on the total profit accumulated from the beginning of the planning horizon up to the current period. Under the same risk measures and when there is a fixed ordering cost, the structure of the optimal policy is unknown.

Finally, we propose a myopic approach based on the CVaR risk measure and we show that this approach, like the exponential utility measure, can preserve (almost) all the structural results as the one under the risk neutral case.

We complement the theoretical results with an experimental study demonstrating that this framework can identify the tradeoff between expected profit and the risk of under-performing. Interestingly, in addition to catering for risk averse inventory managers, we demonstrate empirically that risk averse inventory policies are also powerful when there is a limited demand information. Specifically, we empirically illustrate that when demand information is limited, inventory replenishment policies based on the risk criteria analyzed here yield a profit distribution that dominates those based on risk neutral models.

In Section 2, we review classical and recent approaches in risk averse valuation. In Section 3 we revisited the classical newsvendor problem and derived closed form solutions for a class of quantile based risk averse model. In Section 4, we propose a general framework incorporating risk aversion in multi-period inventory (and pricing) models, and focus on characterizing the optimal inventory policy for a risk averse decision maker. Section 5 presents the computational results illustrating the effects of different risk averse multi-period inventory models on profit distribution. In Section 6 we apply the framework developed in this paper to models with limited demand information. Our objective is to compare empirically risk neutral and risk averse policies in environment when only limited historical data is available. Finally, Section 7 provides some concluding remarks.

We complete this section with a brief statement on notations. Specifically, in this paper, lowercase boldface will be used to denote vectors. A variable with tilde over it such as $\tilde{d}$, denotes a random variable. Similarly, $\tilde{\boldsymbol{d}}$ denotes a vector of random variables.

2 Risk averse valuations

Traditional stochastic inventory models focus on a risk neutral decision maker whose objective is to:

$$\begin{array}{rl} \max & E[f(\mu, \tilde{\boldsymbol{d}})] \\ \text{subject to} & \mu \in \Pi, \end{array} \tag{1}$$

where $\tilde{\boldsymbol{d}}$ denotes a vector of random perturbations of demand, μ is an inventory (and pricing) policy over the set of feasible policies Π, and $f(\mu, \tilde{\boldsymbol{d}})$ is the profit function under evaluation with the control policy, μ and perturbation vector on demand, $\tilde{\boldsymbol{d}}$.

To incorporate the evaluation of risk in inventory models, we recall the notion of Expected Utility Theory (see [22], Chapter 6), which plays a central role in the economics literature and has been applied to model a person's choice under uncertainty. Under the influence of uncertainty, an inventory manager with utility function $u(x)$ optimizes the following problem,

$$\begin{array}{rl} \max & E[u(f(\mu, \tilde{\boldsymbol{d}}))] \\ \text{subject to} & \mu \in \Pi. \end{array} \tag{2}$$

We require the utility function, $u(x)$, to be increasing so that more is always preferred over less. Of course, if $u(x)$ is a linear and increasing function, the model (2) yields the same optimal solution as the risk neutral model of (1). A risk averse decision maker has utility function that is concave so that the marginal satisfaction of gaining a dollar is never more than the marginal loss of satisfaction associated with losing the same amount of money. Hence, it is possible to relate risk aversion to the concavity of the utility function. Specifically, given a utility function, $u(x)$, that is twice differentiable, a well-known measure of risk, the Arrow-Patt coefficient [2, 23] of absolute risk aversion at x, is defined as

$$r_A(x) = -u''(x)/u'(x). \tag{3}$$

A larger magnitude of $r_A(x)$ denotes greater risk aversion at the level x. In particular, the exponential utility function is an important class of utility functions of the form

$$u(x) = 1 - \exp(-x/b), \quad b > 0, \tag{4}$$

for which the Arrow-Patt coefficient, $r_A(x) = 1/b$, is a positive constant. Consequently, a decision maker with exponential utility function has risk aversion that is independent of his/her wealth. A person with greater risk aversion naturally chooses a lower value of the risk parameter b. Interestingly, as we demonstrate later on, the ramification of exponential utility function is the preservation of wealth independent risk averse policies in multi-period inventory models.

In the celebrated mean-variance approach of Markowitz (see Markowitz [20], Huang and Lizenberger [15]), we value an uncertain profit level, $\tilde{z}$, by its mean minus a positive scalar multiple of the variance. In fact the mean-variance valuation of Markowitz satisfies a class of decision makers with concave quadratic utility functions.

Even though the mean-variance approach has been widely used in finance as a risk measure, it has a number of important limitations. Clearly, a concave quadratic function is not always increasing. More importantly, the mean-variance method is intuitively inadequate as it equally penalizes desirable upside influence and undesirable downside outcomes.

Thus, a popular risk measure in finance is Value-at-Risk (VaR) (see Jorion [16], Dowd [11], Duffie and Pan [12]), which allows the decision maker to specify the confidence level for attaining a certain level of wealth. Specifically, in the VaR evaluation of the control policy, μ, one maximizes the quantity, $q_\rho(f(\mu, \tilde{\boldsymbol{d}}))$, where $q_\rho(\tilde{z})$ is the ρ-quantile of a random variable $\tilde{z}$ defined as follows:

$$q_\rho(\tilde{z}) = \inf\{z | \Pr(\tilde{z} \leq z) \geq \rho\}. \tag{5}$$

Unfortunately, even the VaR measure faces some fundamental challenges. As a risk measure, VaR does not preserve the property of subadditivity (see [24, 3]). In other words, a portfolio with two instruments may have a larger risk, i.e., VaR, than the sum of individual VaRs of the two instruments. A different challenge is that the VaR risk measure is indifferent to the extent of which the profit falls below the ρ-quantile level. In particular, given two policies, μ_1 and μ_2, with two corresponding profit distributions, $f(\mu_1, \tilde{\boldsymbol{d}})$ and $f(\mu_2, \tilde{\boldsymbol{d}})$, and with the same ρ-quantile value, the VaR risk evaluation would value both equally, even if the distribution $f(\mu_1, \tilde{\boldsymbol{d}})$ may have greater loss margins below its ρ-quantile profit level. Finally, computing optimal policies under the VaR measure is quite difficult due to the lack of convexity, see [24]. As a result, most approaches for calculating VaR rely on linear approximation of risks and assume normal (or log-normal) distribution of wealth, which can be indeed quite restrictive in inventory models.

In recent years, an alternative measure of risk has evolved under the name of Conditional Value-at-Risk (CVaR), Mean Excess Loss, Mean Shortfall, or Tail VaR. The CVaR criterion (see [24, 25, 3]), which measures the average value of the profit falling below the η-quantile level, has better computational characteristics and surfaces in the financial and insurance literature. With respect to the evaluating of a profit profile, $f(\mu, \tilde{\boldsymbol{d}})$, CVaR ignores the contributions of profit beyond the specified quantile level, and focuses on the average profits from the lower quantiles. For the case of profit maximization, if the underlying distribution has no probability atom at $q_\eta(f(\mu, \tilde{\boldsymbol{d}}))$, the η-CVaR of an inventory (and pricing) policy μ can be defined as follows

$$CVaR_\eta(f(\mu, \tilde{\boldsymbol{d}})) = E[f(\mu, \tilde{\boldsymbol{d}}) | f(\mu, \tilde{\boldsymbol{d}}) \leq q_\eta(f(\mu, \tilde{\boldsymbol{d}}))].$$

When the underlying distribution is discrete, a more general definition of CVaR (see [24, 25]) is as follows:

Definition 2.1

$$CVaR_\eta(f(\mu, \tilde{\boldsymbol{d}})) = \max_{v \in \Re} \left\{ v + \frac{1}{\eta} E[\min(f(\mu, \tilde{\boldsymbol{d}}) - v, 0)] \right\}.$$

Under the CVaR risk assessment parameterized by η, the inventory manager solves the following problem:

$$\begin{array}{rl} \max & CVaR_\eta(f(\mu, \tilde{\boldsymbol{d}})) \\ \text{subject to} & \mu \in \Pi. \end{array} \tag{6}$$

As observed in [25], CVaR is a more consistent measure of risk since it is sub-additive and concave. Moreover, it is easy to observe the following property:

$$CVaR_\eta(f(\mu, \tilde{\boldsymbol{d}})) \leq q_\eta(f(\mu, \tilde{\boldsymbol{d}})).$$

Thus, one can view CVaR risk as a lower bound on the VaR objective. In fact, numerical experiments indicate that the minimization of CVaR also leads to near optimal solutions to models minimizing the VaR measure (see [24]).

More importantly, optimizing CVaR is closely related to second order stochastic dominance defined as follows,

Definition 2.2 *We denote $\tilde{r}_1 \succeq \tilde{r}_2$ to imply that $\tilde{r}_1$ dominates $\tilde{r}_2$ by second order stochastic dominance rule if and only if*

$$E\left[u(\tilde{r}_1)\right] \geq E\left[u(\tilde{r}_2)\right], \text{ for every increasing concave function, } u(x). \tag{7}$$

We denote $\tilde{r}_1 \succ \tilde{r}_2$ for strictly dominates if and only if (7) hold as a strict inequality.

Thus, given two profit returns, $\tilde{r}_1$ and $\tilde{r}_2$, $\tilde{r}_1$ is preferred by all risk averse decision makers if $\tilde{r}_1 \succeq \tilde{r}_2$. Consequently, for any risk averse inventory manager with an increasing concave utility function, optimizing (2) guarantees a profit profile that is not strictly second order dominated by any other feasible profit profile. Interestingly,

Theorem 2.1 *(Levy and Kroll [19]). Let $\tilde{r}_1$ and $\tilde{r}_2$ be two random variables. Then, $\tilde{r}_1 \succeq \tilde{r}_2$ if and only if*

$$E[\tilde{r}_1 \mid \tilde{r}_1 \leq q_\eta(\tilde{r}_1)] \geq E[\tilde{r}_2 \mid \tilde{r}_2 \leq q_\eta(\tilde{r}_2)] \ \forall \eta \in (0,1).$$

The theorem thus implies that as long as the underlying distribution has no probability atom at $q_\eta(\tilde{z})$, optimizing the CVaR inventory model, (6), guarantees a solution with profit distribution that is not dominated by second order dominance rule. These similarities between the CVaR measure and expected utility motivate us to consider both measures in the analysis of risk averse inventory models.

3 The Newsvendor Problem

Consider the celebrated newsvendor model, where a single retailer faces random demand. The retailer places an order to an outside supplier before knowing the level of demand, and sells the product to its customers with a unit price p. Let F be the cumulative distribution function of the demand. The function F is assumed to be strictly increasing and differentiable. For simplicity, we assume that unsatisfied demand is lost and there is no penalty cost for lost sales. In addition, leftover inventory is salvaged with unit price z. Finally, we assume that there is no limit on ordering quantity and that the per unit ordering cost is c with $z < c < p$.

For a given ordering quantity y and any realization of the random demand $\tilde{d}$, the realized profit of the retailer is

$$f(y, \tilde{d}) = -cy + p\min(y, \tilde{d}) + z(y - \tilde{d})^+ = (p - c)y - (p - z)(y - \tilde{d})^+,$$

where $(y - \tilde{d})^+ = \max(y - \tilde{d}, 0)$.

The traditional newsvendor model assumes that the retailer is risk neutral and finds an ordering quantity y^* so as to maximize the retailer expected profit, $E[f(y, \tilde{d})]$. It is well known that the optimal ordering quantity is

$$y^* = F^{-1}\Big(\frac{p - c}{p - z}\Big). \tag{8}$$

Consider now the CVaR risk averse model. Our objective is to find an ordering quantity y_η^* solving the optimization problem

$$\max_y CVaR_\eta(f(y, \tilde{d})).$$

As we demonstrate in Appendix A, there is a simple and intuitive analytical solution for y_η^*.

Theorem 3.1 *Under the CVaR risk averse model parameterized by η, the optimal solution to the newsvendor problem is*

$$y_\eta^* = F^{-1}\Big(\eta\frac{p - c}{p - z}\Big).$$

The above closed form expression suggests that the optimal risk averse ordering quantity is smaller than its risk neutral counterpart. The intuition is clear – since salvage value is smaller than ordering cost, risk is due to overstocking. Thus, the CVaR risk averse model leads to a smaller order quantity. Similar results, under different measures of risk, have appeared in [13] and [6].

Notice the result in Theorem 3.1 can be generalized to the following objective function

$$\max_y \; E[f(y, \tilde{d})] + \xi CVaR_\eta(f(y, \tilde{d})) \; .$$

Details are in Appendix A.

Theorem 3.2 *For any $y > y_1^*$, the profit distribution $f(y, \tilde{d})$ is strictly second order stochastically dominated by the profit distribution $f(y_1^*, \tilde{d})$. That is,*

$$f(y_1^*, \tilde{d}) \succ f(y, \tilde{d}) \quad \forall y > y_1^*$$

Proof : Eeckhoudt et al. [13] showed that for any concave and increasing utility function, $u(\cdot)$, the optimal risk averse ordering quantity, y_u^* is not more than y_1^*. Let

$$g(y) = E[u(f(y, \tilde{d}))],$$

which is a concave function and therefore, we have $g(y_u^*) \geq g(y_1^*) \geq g(y)$ for all $y > y_1^*$. Suppose $g(y) = g(y_1^*)$ for some $y > y_1^*$, by the concavity of the function, $g(\cdot)$, y is also an optimal ordering quantity, which is a contradiction since there does not exist an optimal risk averse ordering quantity that is strictly greater than y_1^* . Hence, for all concave and increasing utility function, $u(\cdot)$ we have

$$E[u(f(y_1^*, \tilde{d}))] > E[u(f(y, \tilde{d}))] \quad \forall y > y_1^*.$$

■

The above observation shows that any risk averse decision maker would prefer profit distributions at production level y_1^* over higher production levels. Hence, to avoid profit profile that is strictly stochastically dominated by second order, the inventory manager would not produce beyond y_1^*.

4 Multi-period Inventory Models

Consider a risk averse firm that has to make replenishment (and pricing) decisions over a finite time horizon with T periods.

Demands in different periods are independent of each other. For each period t, $t = 1, 2, \ldots$, let

$\tilde{d}_t$ = demand in period t
p_t = selling price in period t
$\underline{p}_t, \bar{p}_t$ are lower and upper bounds on p_t, respectively.

Observe that when $\underline{p}_t = \bar{p}_t$ for each period t, price is not a decision variable and the problem is reduced to an inventory control problem. Throughout this paper, we concentrate on demand functions of the following forms:

Assumption 1 *For $t = 1, 2, \ldots$, the demand function satisfies*

$$\tilde{d}_t = D_t(p_t, \tilde{\epsilon}_t) := \tilde{\beta}_t - \tilde{\alpha}_t p_t, \tag{9}$$

where $\tilde{\epsilon}_t = (\tilde{\alpha}_t, \tilde{\beta}_t)$, and $\tilde{\alpha}_t, \tilde{\beta}_t$ are two nonnegative random variables with $E[\tilde{\alpha}_t] > 0$ and $E[\tilde{\beta}_t] > 0$. The random perturbations, $\tilde{\epsilon}_t$, are independent across time.

Let x_t be the inventory level at the beginning of period t, just before placing an order. Similarly, y_t is the inventory level at the beginning of period t after placing an order. The ordering cost function includes both a fixed cost and a variable cost and is calculated for every t, $t = 1, 2, \ldots$, as

$$k\delta(y_t - x_t) + c_t(y_t - x_t),$$

where

$$\delta(x) := \begin{cases} 1, & \text{if } x > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Lead time is assumed to be zero and hence an order placed at the beginning of period t arrives immediately before demand for the period is realized.

Unsatisfied demand is backlogged. Let x be the inventory level carried over from period t to the next period. Since we allow backlogging, x may be positive or negative. A cost $h_t(x)$ is incurred at the end of period t which represents inventory holding cost when $x > 0$ and shortage cost if $x < 0$. For technical reasons, we assume that $h_t(x)$ is convex and $\lim_{|x|\to\infty} h_t(x) = \infty$.

Given a discount factor γ with $0 < \gamma \leq 1$, an initial inventory level, $x_1 = x$, and a replenishment (and pricing) policy, let

$$\tilde{V}_T^{\gamma}(x) = \sum_{t=1}^{T} \gamma^{t-1}(-k\delta(y_t - x_t) - c_t(y_t - x_t) - h_t(x_{t+1}) + p_t D_t(p_t, \tilde{\epsilon}_t)), \tag{10}$$

be the T-period realized discounted profit, where $x_{t+1} = y_t - D_t(p_t, \tilde{\epsilon}_t)$.

The traditional approach for the inventory (and pricing) problem is to decide on ordering (and pricing) policies so as to maximize total expected discounted profit over the entire planning horizon, that is, the objective is to maximize $E[\tilde{V}_T^{\gamma}(x)]$ for any initial inventory level x and any $0 < \gamma \leq 1$.

Denote by $V_t(x)$ the profit-to-go function at the beginning of period t with inventory level x. A natural dynamic program for the risk neutral inventory (and pricing) problem is as follows.

$$V_t(x) = c_t x + \max_{y \geq x, \bar{p}_t \geq p \geq \underline{p}_t} -k\delta(y-x) + g_t(y,p), \tag{11}$$

where $V_{T+1}(x) = 0$ for any x and

$$g_t(y,p) = E\{p(\tilde{\beta}_t - \tilde{\alpha}_t p) - c_t y - h_t(y - (-\tilde{\alpha}_t p + \tilde{\beta}_t)) + \gamma V_{t+1}(y - (-\tilde{\alpha}_t p + \tilde{\beta}_t))\}. \tag{12}$$

Theorem 4.1 *(a) If price is not a decision variable (i.e.,* $\underline{p}_t = \bar{p}_t$ *for each* t*),* $V_t(x)$ *and* $g_t(y,p)$ *are* k*-concave and an* (s,S) *inventory policy is optimal.*

(b) If the demand is additive (i.e., $\tilde{\alpha}_t$ *is a constant),* $V_t(x)$ *and* $\max_{\bar{p}_t \geq p \geq \underline{p}_t} g_t(y,p)$ *are* k*-concave and an* (s,S,p) *policy is optimal.*

(c) For the general case, $V_t(x)$ *and* $g_t(y,p)$ *are symmetric* k*-concave and an* (s,S,A,p) *policy is optimal.*

Part (a) is the classical result proved by Scarf [26] using the concept of k-convexity; Part (b) and Part (c) are proved in Chen and Simchi-Levi [8] using the concepts of k-convexity, for Part (b), and a new concept, the symmetric k-convexity, for Part (c). These concepts are summarized in Appendix B. In fact, the results in [8] hold true under more general demand functions than those in Assumption 1.

In the following subsections, we focus on the inventory and pricing problem, and its special case, the inventory replenishment problem when the objective is to maximize the expected utility or the conditional value-at-risk of the total discounted profit over the entire planning horizon. That is, the objective is to maximize

$$E[u(\tilde{V}_T^\gamma(x))]$$

or

$$CVaR_\eta(\tilde{V}_T^\gamma(x))$$

for any initial inventory level x and any $0 < \gamma \leq 1$.

In the next subsection, we discuss the general framework. This is followed by subsections on the exponential utility and CVaR measures.

4.1 Expected utility and CVaR risk averse models

Unlike the risk neutral models analyzed in Chen and Simchi-Levi [8], the objective function $E[u(\tilde{V}_T^\gamma(x))]$ or $CVaR_\eta(\tilde{V}_T^\gamma(x))$ in its current form appears not to be decomposable and are not amenable to the dynamic programming approach. To deal with this issue, we introduce a new variable w to denote the wealth accumulated from the beginning of the planning horizon up to the current period. Thus, the state of the problem at period t can now be modeled as the inventory level x_t and the accumulated wealth from period 1 to period t, w_t.

Consider the expected utility measure. Let $W_t(x,w)$ be the maximum utility achievable starting at the beginning of period t with an initial inventory level x and an accumulated wealth w. The dynamic program can be written as follows. Let

$$W_{T+1}(x,w) = u(w).$$

and

$$W_t(x,w) = \max_{y \geq x, \bar{p}_t \geq p \geq \underline{p}_t} E_{\tilde{\alpha}_t,\tilde{\beta}_t}[W_{t+1}(x_+, \bar{w}_+)], \tag{13}$$

where $u(\cdot)$ is an increasing concave utility function,

$$x_+ = y - (\tilde{\beta}_t - \tilde{\alpha}_t p)$$

and

$$\bar{w}_+ = w + \gamma^{t-1}(-k\delta(y-x) - c_t(y-x) + p(\tilde{\beta}_t - \tilde{\alpha}_t p) - h_t(y - (\tilde{\beta}_t - \tilde{\alpha}_t p)). \tag{14}$$

Notice that here we assume $W_{T+1}(x,w)$ is independent of x, which implies zero salvage value. Finally, we have

$$\max E[u(\tilde{V}_T^{\gamma}(x))] = W_1(x,0).$$

For the CVaR risk measure, we make a slight modification so that

$$W_{T+1}(x,w) = \min(w,0)$$

while preserving the DP recursion of (13) (for derivation details see Appendix C). Finally, the optimal CVaR can be evaluated as follows:

$$\max CVaR_{\eta}(\tilde{V}_T^{\gamma}(x)) = \max_{v \in \Re} \left\{ v + \frac{1}{\eta} W_1(x,-v) \right\}.$$

Instead of working with the dynamic program (13), we find that it is more convenient to work with an equivalent formulation. Let

$$U_t(x,w) = W_t(x, w - \gamma^{t-1} c_t x).$$

The dynamic program (13) becomes

$$U_t(x,w) = \max_{y \geq x, \bar{p}_t \geq p \geq \underline{p}_t} E_{\tilde{\alpha}_t,\tilde{\beta}_t}[U_{t+1}(x_+, w_+)], \tag{15}$$

where

$$w_+ = w + \gamma^{t-1}(-k\delta(y-x) + f_t(y,p,\tilde{\alpha}_t,\tilde{\beta}_t)),$$

and

$$f_t(y,p,\tilde{\alpha}_t,\tilde{\beta}_t) = -(c_t - \gamma c_{t+1})y + (p - \gamma c_{t+1})(\tilde{\beta}_t - \tilde{\alpha}_t p) - h_t(y - (\tilde{\beta}_t - \tilde{\alpha}_t p)). \tag{16}$$

We have the following observation, which can be easily verified by induction.

Lemma 1 *For any period t and fixed x, $U_t(x,w)$ is increasing in w.*

Interestingly, this observation allows us to show that a wealth dependent base stock inventory policy is optimal in the case $k = 0$ for both the expected utility and the CVaR measures.

Theorem 4.2 *Assume that $k = 0$. In this case, $U_t(x,w)$ is jointly concave in x and w for any period t. Furthermore, a wealth dependent base stock inventory policy is optimal for both the expected utility risk measure and the CVaR risk measure.*

Proof. We prove by induction. Obviously, $U_{T+1}(x, w)$ is jointly concave in x and w. Assume that $U_{t+1}(x, w)$ is jointly concave in x and w. We now prove that a wealth dependent base stock inventory policy is optimal and $U_t(x, w)$ is jointly concave in x and w.

First, notice that for any realization of $\tilde{\alpha}_t$ and $\tilde{\beta}_t$, f_t is jointly concave in (y, p), which implies that w_+ is jointly concave in (w, x, y, p).

Since x_+ is a linear function of (y, p) and w_+ is jointly concave in (w, x, y, p), Lemma 1 allows us to show that $U_{t+1}(x_+, w_+)$ is jointly concave in (w, x, y, p). This implies that $E_{\tilde{\alpha}_t,\tilde{\beta}_t}[U_{t+1}(x_+, w_+)]$ is jointly concave in (w, x, y, p).

We now prove that a $w-$dependent base stock inventory policy is optimal. Let $y^*(w)$ be an optimal solution for the problem

$$\max_y \left\{ \max_{\bar{p}_t \geq p \geq \underline{p}_t} E_{\tilde{\alpha}_t,\tilde{\beta}_t}[U_{t+1}(x_+, w_+)] \right\}.$$

Since $E_{\tilde{\alpha}_t,\tilde{\beta}_t}[U_{t+1}(x_+, w_+)]$ is concave in y for any fixed w, it is optimal to order up to $y^*(w)$ when $x < y^*(w)$ and not to order otherwise. In other words, a state dependent base stock inventory policy is optimal.

Finally, according to Proposition 4 in Appendix B, $U_t(x, w)$ is jointly concave. ■

Recall that in the case of a risk neutral decision maker, a base stock list price policy is optimal. Theorem 4.2 thus implies that in the case of an increasing concave utility risk measure and the CVaR risk measure, the optimal policy is quite different. Indeed, in these cases, the base stock level depends on the total profit accumulated from the beginning of the planning horizon and it is not clear whether a list price policy is optimal. It is also appropriate to point out a difference between optimal policies for a general increasing concave utility risk measure and the CVaR risk measure. Indeed, while the former depends on the initial wealth available at the beginning of the planning horizon, the latter is independent of that amount. This is true since the CVaR measure is invariant to shifting, i.e., $CVaR_\eta(\tilde{X} + C) = CVaR_\eta(\tilde{X}) + C$ for a constant C.

Stronger results exist for models based on the exponential utility risk measure, as is demonstrated in the next subsection.

4.2 Exponential utility risk averse models

We now focus on exponential utility functions of the form $u(w) = 1 - \exp(-w/b)$ with parameter $b > 0$. The beauty of exponential utility functions is that we can essentially separate x and w as is illustrated in the next theorem.

Theorem 4.3 *For any time period* t*, there exists a function* $G_t(x)$ *such that*

$$U_t(x, w) = 1 - \exp(-(w + G_t(x))/b).$$

Proof. We prove by induction. For $t = T + 1$, $G_{T+1}(x) = 0$ for any x. Assume that there exists a function $G_{t+1}(x)$ such that

$$U_{t+1}(x, w) = 1 - \exp(-(w + G_{t+1}(x))/b).$$

From the recursion (13), we have

$$\begin{aligned} U_t(x,w) &= \max_{y\geq x,\bar{p}_t\geq p\geq \underline{p}_t} E_{\tilde{\alpha}_t,\tilde{\beta}_t}[1-\exp(-(w_+ + G_{t+1}(y-\tilde{\beta}_t+\tilde{\alpha}_t p))/b)] \\ &= 1-\exp(-w/b)\min_{y\geq x,\bar{p}_t\geq p\geq \underline{p}_t} \exp(k/b\delta(y-x)-L_t(y,p)/b) \\ &= 1-\exp(-(w+G_t(x))/b), \end{aligned}$$

where

$$L_t(y,p) = -b\ln\Big(E[\exp(-(\gamma^{t-1}f_t(y,p,\tilde{\alpha}_t,\tilde{\beta}_t)+G_{t+1}(y-\tilde{\beta}_t+\tilde{\alpha}_t p))/b)]\Big),$$

and

$$G_t(x) = \max_{y\geq x,\bar{p}_t\geq p\geq \underline{p}_t} -k\delta(y-x)+L_t(y,p). \tag{17}$$

Thus the result is true. ∎

The theorem thus implies that the optimal policy is independent of the accumulated wealth when exponential utility functions are used, which significantly simplifies the problem. In fact, the optimal policy can be found by solving problem (17).

This theorem, together with Theorem 4.2, implies that when $k=0$, a base stock inventory policy is optimal under the exponential utility risk criterion independent of whether price is a decision variable. This is exactly the result obtained by [5] using a more complicated argument for the case when price is not a decision variable. As before, it is not clear whether a list price policy is optimal when $k=0$ and price is a decision variable.

To present our main result for the problem with $k>0$, we need the following theorem.

Theorem 4.4 *If a function f is convex, k-convex or symmetric k-convex, then function*

$$g(x) = \ln(E[\exp(f(x-\tilde{\xi}))])$$

is also convex, k-convex or symmetric k-convex respectively.

Proof. We only prove the case of convexity; the other two cases can be proven by following similar steps.

Define

$$M(x) = E[\exp(f(x-\tilde{\xi}))].$$

and

$$Z(x,w) := E[\exp(w+f(x-\tilde{\xi}))] = e^w M(x).$$

Since f is convex and the exponential function can preserve convexity, $Z(x,w)$ is jointly convex in x and w. In particular, for any x_0, x_1, w_0, w_1 and $\lambda\in[0,1]$, we have from the definition of convexity that

$$Z(x_\lambda, w_\lambda) \leq (1-\lambda)Z(x_0,w_0)+\lambda Z(x_1,w_1), \tag{18}$$

where $x_\lambda = (1-\lambda)x_0+\lambda x_1$ and $w_\lambda=(1-\lambda)w_0+\lambda w_1$.

Dividing both sides of (18) by e^{w_λ}, we have that

$$M(x_\lambda) \leq (1-\lambda)z^\lambda M(x_0) + \lambda z^{-(1-\lambda)}M(x_1),$$

where $z := \exp(w_0 - w_1)$. Taking minimization over z at the righthanded side of the above inequality gives $z = M(x_1)/M(x_0)$. Since w_0 and w_1 are arbitrary, we have that

$$M(x_\lambda) \leq M(x_0)^{1-\lambda} M(x_1)^\lambda.$$

Thus $g(x) = \ln(M(x))$ is convex. ■

We can now present the optimal policy for the risk averse multi-period inventory (and pricing) problem with exponential utility function.

Theorem 4.5 *(a) If price is not a decision variable (i.e., $\underline{p}_t = \bar{p}_t$ for each t), $G_t(x)$ and $L_t(y,p)$ are k-concave and an (s,S) inventory policy is optimal.*

(b) For the general case, $G_t(x)$ and $L_t(y,p)$ are symmetric k-concave and an (s,S,A,p) policy is optimal.

Proof. The main idea is as follows: if $G_{t+1}(x)$ is k-concave when price is not a decision variable (or symmetric k-concave for the general case), then, by Theorem 4.4, $L_t(y,p)$ is k-concave (or symmetric k-concave). The remaining parts follow directly from Lemma 2 and Proposition 2 for k-concavity or Lemma 3 and Proposition 3 for symmetric k-concavity. Lemma 2, Proposition 2, Lemma 3 and Proposition 3 are in Appendix B. ■

We observe the similarities and differences between the optimal policy under the exponential utility measure and the one under the risk neutral case. Indeed, when demand is exogenous, i.e., price is not a decision variable, an (s,S) inventory policy is optimal for the risk neutral case; see Theorem 4.1 Part (a). Theorem 4.5 implies that this is also true under the exponential utility measure. Similarly, for the more general inventory and pricing problem, Theorem 4.1 Part (c) implies that an (s,S,A,p) policy is optimal for the risk neutral case. Interestingly, this policy is also optimal for the exponential utility case.

Of course, the results for the risk neutral case are a bit stronger. Indeed, if demand is additive, Theorem 4.1 Part (b) suggests that an (s,S,p) policy is optimal. Unfortunately, we are not able to prove or disprove such a result for the exponential utility measure.

4.3 Myopic CVaR risk averse models

Unlike the case in which risk is measured based on the expected exponential utility function in Section 4.2, we have to solve a dynamic program with two state variables for the general framework discussed in Section 4.1, which significantly increases computational complexity. Hence, we propose, for the CVaR risk model, a form of myopic risk averse multi-period inventory (and pricing) model for which we are only concerned with the influence of demand uncertainties at the current period. This is done without having to introduce a new variable to represent the accumulated profit, which would increase the dimension and computational complexity of the dynamic program.

Consider the myopic dynamic program

$$V_t^\eta(x) = c_t x + \max_{y \geq x, \bar{p}_t \geq p \geq \underline{p}_t} -k\delta(y-x) + CVaR_{\eta_t}(g_t^\eta(y,p,\tilde{\epsilon}_t)), \tag{19}$$

where $V^{\eta}_{T+1}(x) = 0$ for any x, $\eta = (\eta_1, \ldots, \eta_T)$ with $0 < \eta_t < 1$,

$$g^{\eta}_t(y, p, \tilde{\epsilon}_t) = p(\tilde{\beta}_t - \tilde{\alpha}_t p) - c_t y - h_t(y - (-\tilde{\alpha}_t p + \tilde{\beta}_t)) + \gamma V^{\eta}_{t+1}(y - (-\tilde{\alpha}_t p + \tilde{\beta}_t)). \tag{20}$$

It is easy to see that

$$V^{\eta}_t(x) \leq V_t(x),$$

for any x. In Section 5, we evaluate the effectiveness of this myopic approach.

Before we summarize the structural result for this myopic approach for the risk averse inventory (and pricing) problem, we show that CVaR can preserve k-concavity and symmetric k-concavity.

Proposition 1 *If $g(x, \tilde{\xi})$ is a k-concave (or symmetric k-concave) function of x for any realized $\tilde{\xi}$, then $CVaR_\eta(g(x, \tilde{\xi}))$ is also k-concave (or symmetric k-concave).*

Proof. We prove that $CVaR$ preserves k-concavity. The proof for symmetric k-convexity is similar and hence is omitted. For any x_0, x_1 with $x_0 \leq x_1$ and any $\lambda \in [0, 1]$, let v_0 and v_1 satisfy

$$CVaR_\eta(g(x_0, \tilde{\xi})) = v_0 + \frac{1}{\eta} E_\xi[\min(g(x_0, \tilde{\xi}) - v_0, 0)],$$

and

$$CVaR_\eta(g(x_1, \tilde{\xi} - k)) = v_1 + \frac{1}{\eta} E_\xi[\min(g(x_1, \tilde{\xi}) - k - v_1, 0)].$$

Upon denoting $x_\lambda = (1 - \lambda)x_0 + \lambda x_1$ and $v_\lambda = (1 - \lambda)v_0 + \lambda v_1$, we have that

$$\begin{aligned}
CVaR_\eta(g(x_\lambda, \tilde{\xi})) &\geq v_\lambda + \tfrac{1}{\eta} E_\xi[\min(g(x_\lambda, \tilde{\xi}) - v_\lambda, 0)] \\
&\geq v_\lambda + \tfrac{1}{\eta} E_\xi[\min((1 - \lambda)g(x_0, \tilde{\xi}) + \lambda g(x_1, \tilde{\xi}) - \lambda k - v_\lambda, 0)] \\
&\geq v_\lambda + \tfrac{1}{\eta} E_\xi[(1 - \lambda)\min(g(x_0, \tilde{\xi}) - v_0) + \lambda \min(g(x_1, \tilde{\xi}) - k - v_1, 0)] \\
&= (1 - \lambda)CVaR(g(x_0, \tilde{\xi})) + \lambda CVaR(g(x_1, \tilde{\xi}) - k) \\
&= (1 - \lambda)CVaR(g(x_0, \tilde{\xi})) + \lambda CVaR(g(x_1, \tilde{\xi})) - \lambda k,
\end{aligned}$$

where the first inequality follows from the definition of CVaR, the second inequality from the k-concavity of $g(x, \tilde{\xi})$ for any realized ξ, and the third inequality holds since $\min(x, 0)$ is non-decreasing and concave in x. The last equality follows from the observation that for any constant C,

$$CVaR_\eta(g(x, \tilde{\xi}) - C) = CVaR_\eta(g(x, \tilde{\xi})) - C.$$

Hence $CVaR$ preserves k-convexity. ■

This proposition allows us to prove the structural results for the myopic approach, whose proof is similar to the one for Theorem 4.5 and hence is omitted.

Theorem 4.6 *(a) If price is not a decision variable (i.e., $\underline{p}_t = \bar{p}_t$ for each t), the functions $V^{\eta}_t(x)$ and $CVaR_{\eta_t}(g^{\eta}_t(y, p, \tilde{\epsilon}_t))$ are k-concave. Consequently, an (s, S) inventory policy is optimal.*

(b) For the general case, the functions $V^{\eta}_t(x)$ and $CVaR_{\eta_t}(g^{\eta}_t(y, p, \tilde{\epsilon}_t))$ are symmetric k-concave. Consequently, an (s, S, A, p) policy is optimal.

Discount factor, γ_t	1
Fixed ordering cost, k	100
Unit ordering cost, c_t	1
Unit holding cost, h^+	6
Unit shortage cost, h^-	3
Unit item price, p_t	8

Table 2: Parameters of inventory model

5 Computational Results

In this section, we present the results of an empirical study. The objective is to analyze the effects of different risk averse multi-period inventory models on the profit distributions. In the experimental setup, we consider a fixed price inventory model over a planning horizon with $T = 10$ time periods. The inventory holding and shortage cost function is defined as follows:

$$h_t(y) = h^- \max(-y, 0) + h^+ \max(y, 0),$$

where h^+ is the unit inventory holding cost and h^- is the unit shortage costs. The parameters of the inventory model are listed in Table (2).

Demands in different periods are independent and identically distributed with the following discrete distribution,

$$\tilde{d} = \max(\lfloor 30\tilde{z} \rfloor + 10, 0),$$

where $\tilde{z} \sim \mathcal{N}(0, 1)$, and $\lfloor y \rfloor$, the floor function, denotes the largest integer smaller or equal to y.

To solve the dynamic program associated with the different models there is a need to determine expectations of certain quantities that are difficult to evaluate analytically. For this reason, we apply simulation whenever is necessary. This is done as follows. At each time period, we have a collection of N independent demands, $\tilde{d}^i, i = 1, .., N$, drawn from $\tilde{d}$. Of course, the size N is relatively large compared to T. The inventory policy is derived using this data. Following machine learning terminology, we refer to this data as the *training data.*

Using the training data, we estimate $E[f(\tilde{d})]$,

$$\hat{E}[f(\tilde{d})] = \frac{1}{N} \sum_{i=1}^{N} f(\tilde{d}^i)$$

and construct the policy that maximizes the estimated maximum profit based on the training samples. In all the numerical results presented below, we fix the training set and construct inventory replenishment policies by applying different risk averse criteria.

To evaluate the inventory policies derived from the training set, we analyze the sample profit distributions via Monte Carlo simulation on $S = 10,000$ independent trials. In each trial of evaluating a policy, we generate T demand samples (one for each period) independent from the training set and obtain the accumulated profit at the end of the Tth period.

In addition to analyzing the profit distribution using sample mean and standard deviation, we evaluate the profit distribution using an estimator of $CVaR_\rho(\tilde{y})$, the sample CVaR, which is defined

b	Mean	Standard Deviation	% Negative Profit	Mean Loss
100	259	204.5	8.59	60.4
200	262.8	209.8	9.14	64.9
500	265.9	220.9	10.44	75.2
1000	265.3	227.3	11.54	81
Normal	265.9	234.4	12.26	89.8

Table 3: Summary of results for exponential utility model with $N = 1000$ training data.

as follows

$$\widehat{CVaR}_\rho(\tilde{y}) = \frac{1}{\lfloor \rho S \rfloor} \sum_{j=1}^{\lfloor \rho S \rfloor} \tilde{y}_{(j)},$$

where $\tilde{y}_{(k)}$ is the kth order statistics such that $\tilde{y}_{(1)} \leq ... \leq \tilde{y}_{(S)}$. Given two sets of S independent samples of random variables $v(\omega)$ and $u(\omega)$, if S is reasonably large and

$$\widehat{CVaR}_\rho(v(\omega)) > \widehat{CVaR}_\rho(u(\omega)) \quad \forall \rho \in [1/S, 1),$$

then we expect that $v(\tilde{\omega})$ is likely to dominate $u(\tilde{\omega})$ by second order stochastic dominance, that is, $v(\tilde{\omega})$ is strictly preferred over $u(\tilde{\omega})$ by every investors with increasing concave utility functions.

In the experimental setup, we focus on two risk averse models, namely the Exponential Utility Inventory Model and the Myopic CVaR Inventory Model.

5.1 Exponential Utility Approach

We consider a utility function of the form $u(w) = 1 - \exp(-w/b)$, where b is chosen from the set $\{100, 200, 500, 1000\}$. We fix the training set to the size of $N = 1000$. For each risk parameter b we construct the inventory policy using the training set. Observe that for small value of b, for instance, $b = 20$, $u(w)$ can be extremely large, which is generally beyond the precision of floating point representation.

Table (3) shows the profit profile as the parameter b varies. The column under '% Negative Profit' indicates the fraction of profit samples falling below zero. Likewise, 'Mean Loss' refers to the average loss, given that the profit is negative. The row 'Normal' represents the risk neutral case.

Observe that as the parameter b decreases, the sample mean profit and standard deviation of the profit profile decrease, which indicates greater risk averse decisions.

We compare the different policies, generated using different values of b, using the sample CVaR. Figure (1) presents the sample CVaR as a function of ρ for each b. Observe that the curves cross, which suggests that the profit distributions generated by varying the parameter b do not appear to be second order stochastically dominated by each other. This is not surprising, as indicated by Theorem 2.1, since $u(w)$ is an increasing concave utility function.

5.2 Myopic CVaR Approach

We now experiment with the myopic CVaR inventory model of (19), for which a wealth independent (s, S) policy is optimal. Of course, being a myopic approach, we cannot guarantee that the policy

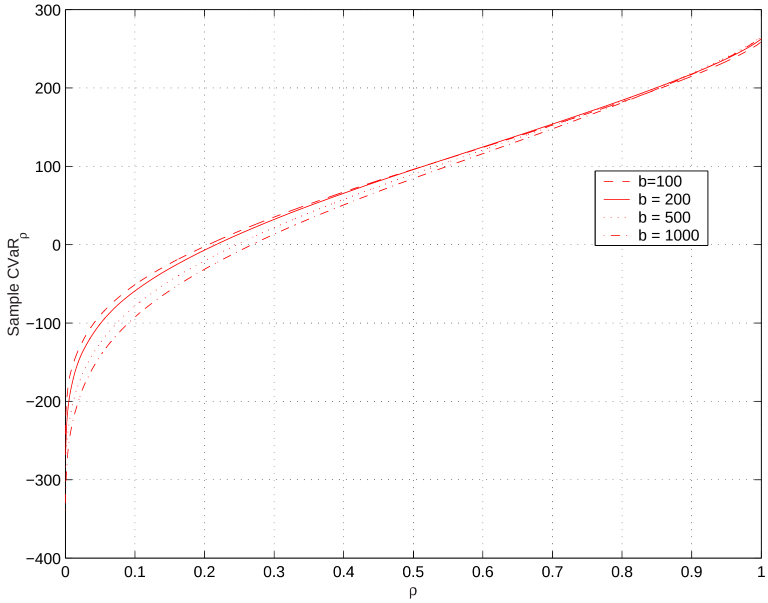

Figure 1: Comparison of sample CVaR for policies derived from $N = 1000$ samples using exponential utility model.

derived from this method would lead to profit profile that is not second order stochastically dominated. Nevertheless, unlike the exponential utility approach, we can potentially vary the parameters for risk adjustment, η_t, without causing any computational issues. We fix the parameters $\eta_t = \eta$ for every time period and analyze the profit profile as η varies in the set $\{0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$.

The results with $N = 1000$ training samples are presented in Table (4). We observe that the sample mean and standard deviation behave similarly to their behavior in the case of exponential utility, that is, as the parameter η decreases from 1, which is the risk neutral level (denoted by 'Normal'), the sample mean and standard deviation decrease. We also observe the trend that as η decreases, i.e., as the manager becomes more risk averse, the average loss, and the likelihood of a negative profit, decrease. Finally, Figure (2) suggests that the profit distributions generated by varying the parameter η do not appear to be second order stochastically dominated by each other.

6 Impact of Limited Demand Information

An important challenge faced by most traditional stochastic inventory models is that they require a complete knowledge of the demand distributions, which is unrealistic in many practical situations. Evidently, inaccurate estimates of demand distributions yield inappropriate replenishment and pricing decisions, leading to poor inventory policies; see computational results in [4]. Hence, given limited knowledge of demand distributions, there is an implicit "risk" of obtaining poor inventory policies.

We assume that while the decision maker has no precise information of demand distributions, she has N observations, the training sample, for each time period. These observations, representing historical data, have been generated independently from an identical distribution for that period. A natural approach for a risk neutral decision maker is to use this data to estimate the demand

η	Mean	Standard Deviation	% Negative Profit	Mean Loss
0.7	245.9	197.3	9.06	58.1
0.75	253.3	200.4	8.77	58.4
0.8	260.4	209.2	9.26	64.4
0.85	263.5	215.3	9.93	70.2
0.9	265.9	220.9	10.44	75.2
0.95	265.3	227.3	11.54	81
Normal	265.9	234.4	12.26	89.8

Table 4: Summary of results for myopic CVaR model with $N = 1000$ training data.

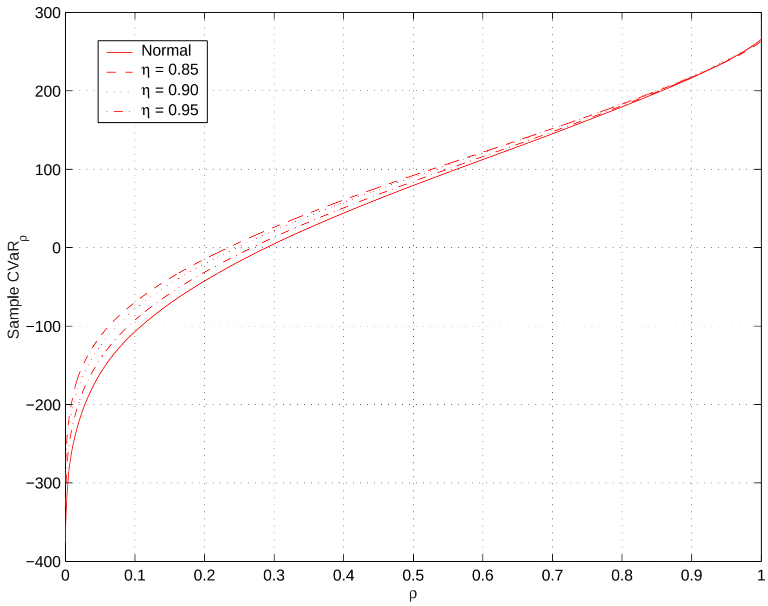

Figure 2: Comparison of sample CVaR for policies derived from $N = 1000$ samples using Myopic CVaR model.

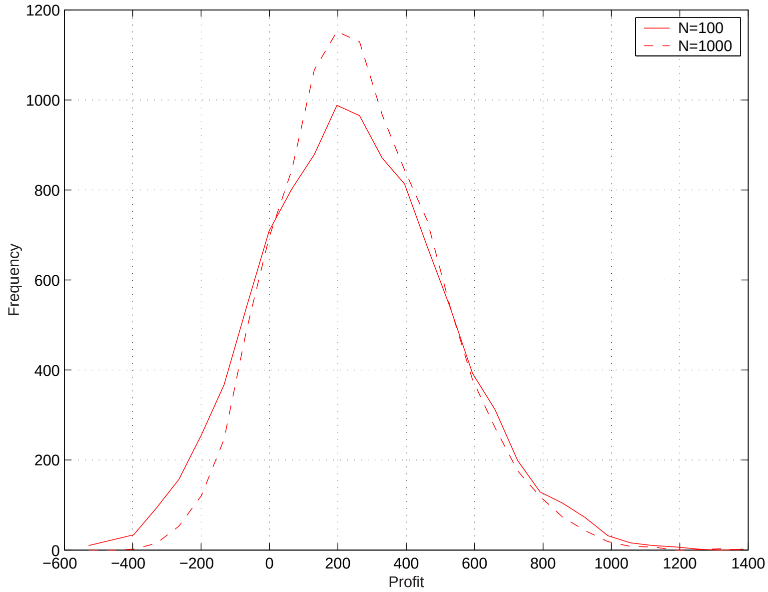

Figure 3: Comparison of distributions for risk neutral policies derived from $N = 100$ and $N = 1000$ samples.

distribution and hence the appropriate replenishment policy.

Clearly, the larger the sample size N, the closer the policy to the optimal policy. However, working with limited demand samples, the policies could significantly differ from the optimal policy. In the following experiment, we construct inventory policies using $N = 100$ instead of $N = 1000$, as was done in the previous section. The small size N represents the more realistic situation in which one has limited knowledge of the demand distributions.

Figure (3) illustrates the difference in profit distributions between the risk neutral policies derived using $N = 1000$ and $N = 100$ training samples. With $N = 100$ training samples, the sample mean profit and standard deviation are respectively 250.1 and 272.4, while with $N = 1000$ training samples, the mean and standard deviation improve (that is, higher mean and lower standard deviation) to respectively, 265.9 and 234.4.

In Figure (4), we compare the respective sample CVaR of the profit distributions derived with $N = 100$ and $N = 1000$ training data. That is, for any ρ we determine the sample CVaR for the above profit distributions, those associated with risk neutral policies based on $N = 100$ and $N = 1000$. It is very suggestive that the solutions derived from only $N = 100$ training data is second order stochastically dominated by the solution with $N = 1000$ training data. Thus, a policy based on the risk neutral model and limited demand information may lead to poor profit profile.

We therefore study the profit distribution obtained when the policy is derived using limited data and the risk averse inventory models. As we shall demonstrate empirically, these profit distributions generated based on the CVaR and exponential measures dominate the profit distribution obtained using the risk neutral model.

In Table (5) we compared policies based on the exponential utility approach and policies based on the risk neutral approach, in both cases using $N = 100$ training samples. Observe that unlike

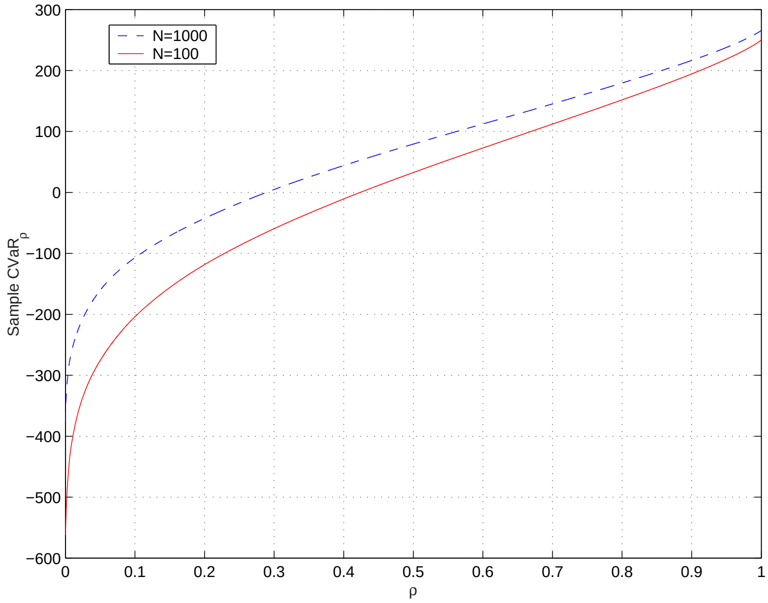

Figure 4: Comparison of sample CVaR for risk neutral policies derived from $N = 100$ and $N = 1000$ samples.

b	Mean	Standard Deviation	% Negative Profit	Mean Loss
100	263.5	215.3	9.93	70.2
200	265.3	227.3	11.54	81
500	264	241.2	13.38	96.8
1000	261.4	248.6	14.52	104.6
Normal	250.1	272.4	18.02	132.4

Table 5: Summary of results for exponential utility model with $N = 100$ training data.

the case of larger training samples, the performance of the risk neutral inventory policy ('Normal') is significantly worse than all the risk averse policies. In particular, with the parameter $b = 200$, the sample mean is significantly improved by 6.1%, the standard deviation reduced by 16.6% and the mean loss is reduced by 38.8%.

Figure (5) presents the sample CVaR of each one of the four policies; the policy based on the risk neutral model and those based on the exponential utility approach with $b = 200, 500, 1000$. The figure suggests that the profit profile based on a risk neutral model is second order stochastically dominated by the profit profiles associated with the exponential utility approach.

In Table (6) we present a similar experiment using the myopic CVaR approach with $N = 100$ training data. We observe similar trends, that is, the profit mean (standard deviation) of the myopic CVaR policy is larger (smaller) than the risk neutral policy for all values of η considered. In particular, when the parameter $\eta = 0.85$, the sample mean improved by 5% and the mean loss level is reduced by 27% over the risk neutral approach. Similarly to the exponential utility approach, Figure (6) suggests

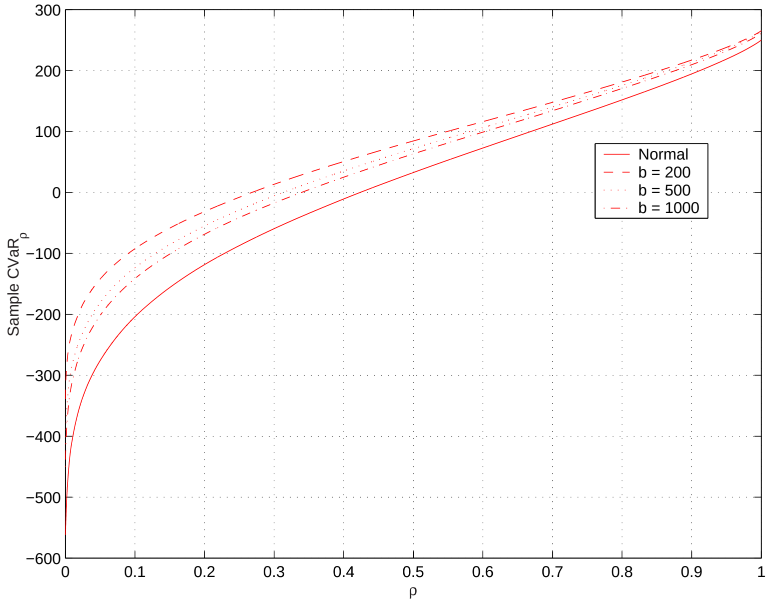

Figure 5: Comparison of sample CVaR for policies derived from $N = 100$ samples using exponential utility model.

that the profit profile based on a risk neutral model is second order stochastically dominated by the profit profiles associated with the three policies derived using the myopic CVaR model.

Hence, this empirical study suggests that when little is known about the demand distributions, the inventory policies derived from a risk neutral model could lead to poor profit performance. More importantly, our preliminary computational study suggests that policies based on the risk averse models can lead to stochastically dominating profit profiles, underscoring the importance of these risk averse approaches in practical application of inventory models.

η	Mean	Standard Deviation	% Negative Profit	Mean Loss
0.7	260.6	226.8	12.02	81.1
0.75	262.1	233.8	12.58	89.7
0.8	260.3	240.9	13.74	97.2
0.85	262.5	241	13.59	96.5
0.9	259.9	248.5	14.71	104.8
0.95	254.4	264	16.55	124.3
Normal	250.1	272.4	18.02	132.4

Table 6: Summary of results for myopic CVaR model with $N = 100$ training data.

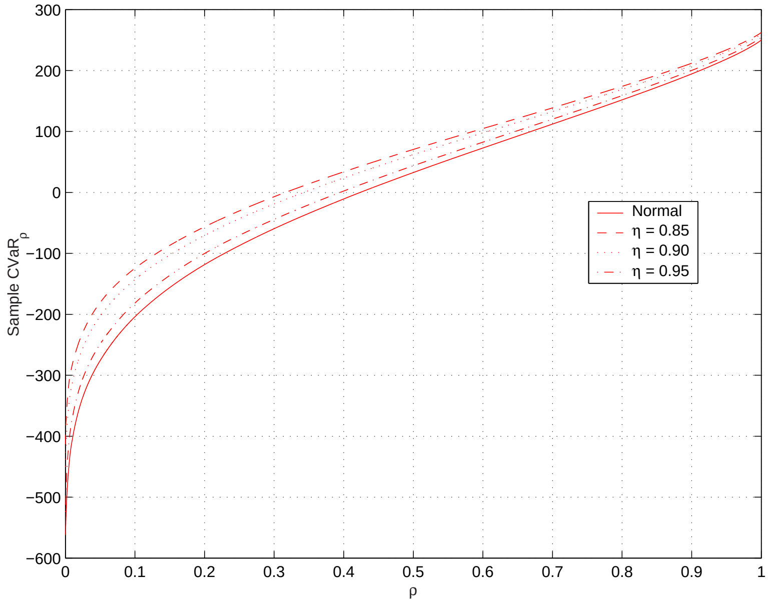

Figure 6: Comparison of sample CVaR for myopic CVaR model.

7 Conclusions

In this paper, we propose a general framework to incorporate risk aversion into inventory (and pricing) models. The framework proposed in this paper and the results obtained motivate a number of extensions.

- **Risk Averse Infinite Horizon Models:** The risk averse infinite horizon models are not only important but also theoretically challenging. The analysis of these types of models is presented in a companion paper.

- **The Stochastic Cash Balance Problem:** Recently, Chen and Simchi-Levi [10] applied the concept of symmetric k-convexity and its extension to characterize the optimal policy for the classical stochastic cash balance problem when the decision maker is risk neutral. It turns out, similar to what we did in Section 4.2, that this type of policies remains optimal for risk averse cash balance problems under exponential utility measure.

- **Random Yield Models:** So far we have assumed that uncertainty is only associated with the demand process. An important challenge is to incorporate supply uncertainty into these risk averse inventory problems.

- **Capacity Constraints:** When there is no fixed ordering cost, capacity constraints on inventory levels can be incorporated into our risk averse inventory models. In this case, one can easily verify that a wealth dependent modified bask stock inventory policy is optimal when risk is evaluated by a general increasing concave utility function or by the CVaR measure. On the other hand, a wealth independent modified base stock inventory policy is optimal when risk is measured by the exponential utility function. In addition, this result allows us to incorporate spot market and portfolio contracts into our risk averse multi-period framework. Observe that, a different risk averse model, based on the mean-variance tradeoff in supply contracts, cannot be easily extended to multi-period, as pointed out by Martínez-de-Albéniz and Simchi-Levi [21].

Of course, it is also interesting to extend the framework proposed in this paper to more general inventory models, such as the multi-echelon inventory models and continuous time inventory (and pricing) models. Finally, it is possible to extend this framework to different environments, those that go beyond inventory models, for instance, revenue management models.

A Proof of Theorem 3.1 and its generalization

We first proof Theorem 3.1. Then we present its generalization.
Proof. Define a concave function

$$g(y,v) := v + \frac{1}{\eta}\int_D [(p-c)y - (p-z)(y-D)^+ - v]^- dF(D).$$

According to Proposition 2.1, the optimal ordering quantity $y^* = \arg\max_{y\geq 0}\{\max_v g(y,v)\}$. Obviously

$$g(y,v) = v + \frac{1}{\eta}\left(\int_0^y [(z-c)y + (p-z)D - v]^- dF(D) + \int_y^\infty [(p-c)y - v]^- dF(D)\right).$$

For any fixed y, we distinguish between three different cases:

(1) $v < (z-c)y$.
In this case, $g(y,v) = v$ and thus $\frac{\partial g}{\partial v} = 1 > 0$.

(2) $(z-c)y \leq v \leq (p-c)y$
In this case, one can derive that

$$g(y,v) = v + \frac{1}{\eta}\left\{[(z-c)y - v]F(\frac{v-(z-c)y}{p-z}) + (p-z)\int_0^{\frac{v-(z-c)y}{p-z}} DdF(D)\right\},$$

and

$$\frac{\partial g}{\partial v} = 1 - \frac{F(\frac{v-(z-c)y}{p-z})}{\eta}.$$

(3) $v > (p-c)y$
In this case, we have that

$$g(y,v) = v + \frac{1}{\eta}\left\{(z-p)yF(y) + (p-z)\int_0^y Df(D)dD + (p-c)y + v\right\},$$

and

$$\frac{\partial g}{\partial v} = \frac{\eta - 1}{\eta} < 0.$$

It is clear that, for a fixed y, $g(y,z)$ attains maximum when $(z-c)y \le v \le (p-c)y$. Let $v^*(y)$ be an optimal solution of $g(y,z)$ for a fixed y. If $y \ge F^{-1}(\eta)$, we have $v^*(y) = (p-z)F^{-1}(\eta) + (z-c)y$; otherwise, $v^*(y) = (p-c)y$.

When $y \ge F^{-1}(\eta)$, we have that

$$g(y, v^*(y)) = (z-c)y + \frac{p-z}{\eta}\int_0^\eta F(x)dx,$$

and

$$\frac{dg}{dy} = z - c < 0.$$

On the other hand, when $y < F^{-1}(\eta)$,

$$g(y, v^*(y)) = (p-c)y + \frac{1}{\eta}\left\{(z-p)yF(y) + (p-z)\int_0^y DdF(D)\right\},$$

and

$$\frac{dg}{dy} = (p-c) + \frac{1}{\eta}(z-p)F(y).$$

Therefore, $y^* = F^{-1}\left(\eta\frac{p-c}{p-z}\right)$. ■

Now we consider the objective function that trades off between the expected profit and CVaR.

Theorem A.1 *Under the following CVaR risk averse model*

$$\max_y \; E[f(y,\tilde{d})] + \xi CVaR_\eta(f(y,\tilde{d})) \ ,$$

the optimal newsvendor solution is

$$y^* = \begin{cases} F^{-1}\left(\frac{p-c}{p-z} - \frac{\xi(c-z)}{p-z}\right) & \textit{when} \;\; \eta \le \frac{p-c}{p-z} - \frac{\xi(c-z)}{p-z} \\ F^{-1}\left(\frac{\eta\xi+\eta}{\eta+\xi}\frac{p-c}{p-z}\right) & \textit{when} \;\; \eta > \frac{p-c}{p-z} - \frac{c-z}{\xi(p-z)} \end{cases} .$$

Notice from the above theorem we can directly observe the following special cases corresponding to the risk neutral and pure CVaR objective models

1. When $\xi = 0$, the optimal ordering quantity is the risk neutral decision $y^* = F^{-1}\left(\frac{p-c}{p-z}\right)$.

2. When $\xi \to \infty$ we have $y^* = F^{-1}\left(\eta\frac{p-c}{p-z}\right)$

Proof. Define a concave function

$$\begin{aligned} g(y,v) \;\; &:= \;\; \xi\left(v + \frac{1}{\eta}\int_D [(p-c)y - (p-z)(y-D)^+ - b(D-y)^+ - v]^- dF(D)\right) \\ &\quad + \int_D [(p-c)y - (p-z)(y-D)^+ - b(D-y)^+]dF(D) \ . \end{aligned}$$

According to Proposition 2.1, the optimal ordering quantity $y^* = \arg\max_{y\ge 0}\{\max_v g(y,v)\}$.

Obviously

$$\begin{aligned}
g(y,v) &= \xi\left[v+\frac{1}{\eta}\left(\int_0^y[(z-c)y+(p-z)D-v]^- dF(D)+\int_y^\infty[(p-c+b)y-bD-v]^- dF(D)\right)\right] \\
&\quad +\left(\int_0^y[(z-c)y+(p-z)D]dF(D)+\int_y^\infty[(p-c+b)y-bD]dF(D)\right) .
\end{aligned}$$

For any fixed y, we distinguish between three different cases:

(1) $v<(z-c)y$.

In this case,

$$g(y,v)=\xi v+\left(\int_0^y[(z-c)y+(p-z)D]dF(D)+\int_y^\infty[(p-c+b)y-bD]dF(D)\right) ,$$

thus $\frac{\partial g}{\partial v}=\xi>0$.

(2) $(z-c)y\le v\le(p-c)y$

In this case, one can derive that

$$\begin{aligned}
g(y,v) &= \xi\Big(v+\frac{1}{\eta}\Big\{[(z-c)y-v]F\left(\frac{v-(z-c)y}{p-z}\right)+(p-z)\int_0^{\frac{v-(z-c)y}{p-z}}DdF(D) \\
&\quad [(p+b-c)y-v]\left[F\left(\frac{(p+b-c)y-v}{b}\right)-F(y)\right]-b\int_y^{\frac{(p+b-c)y-v}{b}}DdF(D)\Big\}\Big) \\
&\quad +\left(\int_0^y[(z-c)y+(p-z)D]dF(D)+\int_y^\infty[(p-c+b)y-bD]dF(D)\right) ,
\end{aligned}$$

and

$$\begin{aligned}
\frac{\partial g}{\partial v} &= \xi-\frac{\xi}{\eta}\left[F\left(\frac{v-(z-c)y}{p-z}\right)-F\left(\frac{(p+b-c)y-v}{b}\right)+F(y)\right] \\
\frac{\partial g}{\partial y} &= \frac{\xi}{\eta}\Big((z-c)F\left(\frac{v-(z-c)y}{p-z}\right)+(p+b-c)F\left(\frac{(p+b-c)y-v}{b}\right) \\
&\quad -(p+b-c)F(y)-[(p-c)y-v]f(y)\Big)+(p+b-c)-(p+b-z)F(y)
\end{aligned}$$

Consider the following two cases:

a) $y<F^{-1}(\eta)$

$$\begin{aligned}
\frac{\partial g}{\partial v} &= \xi\left\{1-\frac{1}{\eta}\left[F\left(\frac{v-(z-c)y}{p-z}\right)-F\left(\frac{(p+b-c)y-v}{b}\right)+F(y)\right]\right\} \\
&> \xi\left\{1-\frac{1}{\eta}\left[F\left(\frac{v-(z-c)y}{p-z}\right)-F\left(\frac{(p+b-c)y-v}{b}\right)+\eta\right]\right\}>0
\end{aligned}$$

The last inequality follows from $\frac{v-(z-c)y}{p-z}>\frac{(p+b-c)y-v}{b}$ because $v<(p-c)y$.

b) $y>F^{-1}(\eta)$

(3) $(p-c)y < v \le (p+b-c)y$

In this case, one can derive that

$$\begin{aligned} g(y,v) &= \xi\left(v+\frac{1}{\eta}\left\{[(z-c)y-v]F\left(\frac{v-(z-c)y}{p-z}\right)+(p-z)\int_0^{\frac{v-(z-c)y}{p-z}} DdF(D)\right\}\right) \\ &\quad +\left(\int_0^y [(z-c)y+(p-z)D]dF(D)+\int_y^\infty [(p-c+b)y-bD]dF(D)\right) , \end{aligned}$$

and

$$\frac{\partial g}{\partial v} = \xi\left(1-\frac{F(\frac{v-(z-c)y}{p-z})}{\eta}\right) .$$

(4) $v > (p+b-c)y$

In this case, we have that

$$\begin{aligned} g(y,v) &= \xi\Big(v\left(1-\frac{1}{\eta}\right)+\frac{1}{\eta}\cdot \\ &\quad \Big\{(z-p-b)yF(y)+(p-z)\int_0^y DdF(D)+(p-c+b)y-b\int_y^\infty DdF(D)\Big\}\Big) \\ &\quad +\left(\int_0^y [(z-c)y+(p-z)D]dF(D)+\int_y^\infty [(p+b-c)y-bD]dF(D)\right) , \end{aligned}$$

and

$$\frac{\partial g}{\partial v} = \xi\frac{\eta-1}{\eta} < 0.$$

It is clear that, for a fixed y, $g(y,z)$ attains maximum when $(z-c)y \le v \le (p+b-c)y$. Let $v^*(y)$ be an optimal solution of $g(y,z)$ for a fixed y. We have the following cases for y.

a)
$$\frac{p-z}{p+b-z}F^{-1}(\eta) \le y \le F^{-1}(\eta),$$

we have
$$v^*(y) = (p-z)F^{-1}(\eta)+(z-c)y ,$$

which satisfies $(p-c)y < v(y^*) \le (p+b-c)y$.

In this case, we have that

$$\begin{aligned} g(y,v^*(y)) &= \xi\left((z-c)y+\frac{p-z}{\eta}\int_0^\eta F^{-1}(x)dx\right) \\ &\quad +\left(\int_0^y [(z-c)y+(p-z)D]dF(D)+\int_y^\infty [(p+b-c)y-bD]dF(D)\right) , \end{aligned}$$

and
$$\frac{dg}{dy} = \xi(z-c)+(p+b-c)+(z-p-b)F(y) .$$

Thus when

$$\frac{(p+b-c)-\xi(c-z)}{p+b-z} \leq \eta \leq F\left(\frac{p+b-z}{p-z}F^{-1}\left(\frac{\xi(z-c)+(p+b-c)}{p+b-z}\right)\right) ,$$

and

$$\xi < \frac{p+b-c}{c-z}$$

we have

$$y^* = F^{-1}\left(\frac{(p+b-c)-\xi(c-z)}{p+b-z}\right)$$

b)

$$y < \frac{p-z}{p+b-z}F^{-1}(\eta) ,$$

we have $v^*(y) = (p+b-c)y$.

$$\begin{aligned} g(y, v^*(y)) &= \xi\left((p+b-c)y + \frac{1}{\eta}\left\{(z-p-b)yF(y) + (p-z)\int_0^y DdF(D) - b\int_y^\infty DdF(D)\right\}\right) \\ &+ \left(\int_0^y [(z-c)y + (p-z)D]dF(D) + \int_y^\infty [(p+b-c)y - bD]dF(D)\right) , \end{aligned}$$

and

$$\frac{dg}{dy} = (1+\xi)(p+b-c) + (1+\frac{\xi}{\eta})(z-p-b)F(y) .$$

Thus when

$$\xi < \frac{\eta\left(\frac{p+b-z}{p+b-c}F\left(\frac{p-z}{p+b-z}F^{-1}(\eta)\right)-1\right)}{\eta - \frac{p+b-z}{p+b-c}F\left(\frac{p-z}{p+b-z}F^{-1}(\eta)\right)}$$

we have

$$y^* = F^{-1}\left(\frac{\eta\xi+\eta}{\eta+\xi}\frac{p+b-c}{p+b-z}\right)$$

c) otherwise,

$$v^*(y) = (p+b-c)y .$$

When $y \geq \frac{p-z}{p+b-z}F^{-1}(\eta)$,
On the other hand, when $y < F^{-1}(\eta)$,

$$\begin{aligned} g(y, v^*(y)) &= \xi\left((p-c)y + \frac{1}{\eta}\left\{(z-p)yF(y) + (p-z)\int_0^y DdF(D)\right\}\right) \\ &+ \left(\int_0^y [(z-c)y + (p-z)D]dF(D) + \int_y^\infty (p-c)ydF(D)\right) , \end{aligned}$$

and

$$\frac{dg}{dy} = (1+\xi)(p-c) + (1+\frac{\xi}{\eta})(z-p)F(y) .$$

Therefore, for $\xi > 0$, we have

$$y^* = \begin{cases} F^{-1}\left(\frac{p-c}{p-z} - \frac{\xi(c-z)}{p-z}\right) & \text{when } \eta \le \frac{p-c}{p-z} - \frac{\xi(c-z)}{p-z} \\ F^{-1}\left(\frac{\eta\xi+\eta}{\eta+\xi}\frac{p-c}{p-z}\right) & \text{when } \eta > \frac{p-c}{p-z} - \frac{\xi(c-z)}{p-z} \end{cases}.$$

■

B Review on k-convexity and symmetric k-convexity

In this section, we review some important properties of k-convexity and symmetric k-convexity that are used in this paper; see Chen [7] for more details.

The concept of k-convexity was introduced by Scarf [26] to prove the optimality of an (s, S) for the traditional inventory control problem.

Definition B.1 *A real-valued function f is called k-convex for $k \ge 0$, if for any $x_0 \le x_1$ and $\lambda \in [0, 1]$,*

$$f((1-\lambda)x_0 + \lambda x_1) \le (1-\lambda)f(x_0) + \lambda f(x_1) + \lambda k. \tag{21}$$

Below we summarize properties of k-convex functions.

Lemma 2 *(a) A real-valued convex function is also 0-convex and hence k-convex for all $k \ge 0$. A k_1-convex function is also a k_2-convex function for $k_1 \le k_2$.*

(b) If $f_1(y)$ and $f_2(y)$ are k_1-convex and k_2-convex respectively, then for $\alpha, \beta \ge 0$, $\alpha f_1(y) + \beta f_2(y)$ is $(\alpha k_1 + \beta k_2)$-convex.

(c) If $f(y)$ is k-convex and w is a random variable, then $E\{f(y-w)\}$ is also k-convex, provided $E\{|f(y-w)|\} < \infty$ for all y.

(d) Assume that f is a continuous k-convex function and $f(y) \to \infty$ as $|y| \to \infty$. Let S be a minimum point of g and s be any element of the set

$$\{x | x \le S, f(x) = gf(S) + k\}.$$

Then the following results hold.

(i) $f(S) + k = f(s) \le f(y)$, for all $y \le s$.

(ii) $f(y)$ is a non-increasing function on $(-\infty, s)$.

(iii) $f(y) \le f(z) + k$ for all y, z with $s \le y \le z$.

Proposition 2 *If $f(x)$ is a K-convex function, then function*

$$g(x) = \min_{y \ge x} Q\delta(y - x) + f(y),$$

is $\max\{K, Q\}$-convex.

Recently a weaker concept of symmetric k-convexity was introduced by Chen and Simchi-Levi [8, 9] when they analyze the joint inventory and pricing problem with fixed ordering cost.

Definition B.2 *A function* $f : \Re \to \Re$ *is called symmetric* k*-convex for* $k \geq 0$*, if for any* $x_0, x_1 \in \Re$ *and* $\lambda \in [0, 1]$,

$$f((1-\lambda)x_0 + \lambda x_1) \leq (1-\lambda)f(x_0) + \lambda f(x_1) + \max\{\lambda, 1-\lambda\}k. \tag{22}$$

A function f *is called symmetric* k*-concave if* $-f$ *is symmetric* k*-convex.*

Observe that k-convexity is a special case of symmetric k-convexity. The following results describe properties of symmetric k-convex functions, properties that are parallel to those summarized in Lemma 2 and Proposition 2. Finally, notice that the concept of symmetric k-convexity can be easily extended to *multi-dimensional* space.

Lemma 3 *(a) A real-valued convex function is also symmetric* 0*-convex and hence symmetric* k*-convex for all* $k \geq 0$*. A symmetric* k_1*-convex function is also a symmetric* k_2*-convex function for* $k_1 \leq k_2$*.*

(b) If $g_1(y)$ *and* $g_2(y)$ *are symmetric* k_1*-convex and symmetric* k_2*-convex respectively, then for* $\alpha, \beta \geq 0$, $\alpha g_1(y) + \beta g_2(y)$ *is symmetric* $(\alpha k_1 + \beta k_2)$*-convex.*

(c) If $g(y)$ *is symmetric* k*-convex and* w *is a random variable, then* $E\{g(y - w)\}$ *is also symmetric* k*-convex, provided* $E\{|g(y - w)|\} < \infty$ *for all* y*.*

(d) Assume that g *is a continuous symmetric* k*-convex function and* $g(y) \to \infty$ *as* $|y| \to \infty$*. Let* S *be a global minimizer of* g *and* s *be any element from the set*

$$X := \{x | x \leq S, g(x) = g(S) + k \text{ and } g(x') \geq g(x) \text{ for any } x' \leq x\}.$$

Then we have the following results.

(i) $g(s) = g(S) + k$ *and* $g(y) \geq g(s)$ *for all* $y \leq s$*.*

(ii) $g(y) \leq g(z) + k$ *for all* y, z *with* $(s + S)/2 \leq y \leq z$*.*

Proposition 3 *If* $f(x)$ *is a symmetric* K*-convex function, then the function*

$$g(x) = \min_{y \leq x} Q\delta(x - y) + f(y)$$

is symmetric $\max\{K, Q\}$*-convex. Similarly, the function*

$$h(x) = \min_{y \geq x} Q\delta(x - y) + f(y)$$

is also symmetric $\max\{K, Q\}$*-convex.*

Proposition 4 *Let* $f(\cdot, \cdot)$ *be a function defined on* $\Re^n \times \Re^m \to \Re$*. Assume that for any* $x \in \Re^n$ *there is a corresponding set* $C(x) \subset \Re^m$ *such that the set* $C \equiv \{(x, y) \mid y \in C(x), x \in \Re^n\}$ *is convex in* $\Re^n \times \Re^m$*. If* f *is symmetric* k*-convex over the set* C*, and the function*

$$g(x) = \inf_{y \in C(x)} f(x, y)$$

is well defined, then g *is symmetric* k*-convex over* $\Re^n$*.*

Proof. For any $x_0, x_1 \in \Re^n$ and $\lambda \in [0,1]$, let $y_0, y_1 \in \Re^m$ such that $g(x_0) = f(x_0, y_0)$ and $g(x_1) = f(x_1, y_1)$. Then

$$\begin{aligned} g((1-\lambda)x_0 + \lambda x_1) &\leq f((1-\lambda)x_0 + \lambda x_1, (1-\lambda)y_0 + \lambda y_1) \\ &\leq (1-\lambda)f(x_0, y_0) + \lambda f(x_1, y_1) + \max\{\lambda, 1-\lambda\}K \\ &= (1-\lambda)g(x_0) + \lambda g(x_1) + \max\{\lambda, 1-\lambda\}K, \end{aligned}$$

Therefore g is symmetric K-convex. ■

C Derivation of CVaR DP formulation

Referring to Proposition 2.1, we have

$$\max_{y_t, p_t: 1 \leq t \leq T} (CVaR_\eta(\tilde{V}_T^\gamma(x))) = \max_v G_x(v) \ ,$$

where

$$G_x(v) := v + \frac{1}{\eta} \max_{y_t, p_t: 1 \leq t \leq T} \{E_{\tilde{\alpha}_t, \tilde{\beta}_t: 0 \leq t \leq T}[(\tilde{V}_T^\gamma(x) - v)^-]\}.$$

We can design the following algorithm: for any fixed v, we solve $G_x(v)$ using dynamic program. Then find the optimal v that maximizes $G_x(v)$, which leads to the solution to the CVaR maximization problem.

The DP algorithm for solving $G_x(v)$ is defined as

$$\bar{W}_{T+1}^v(x, w) = v + \frac{1}{\eta}(w - v)^-$$

and

$$\bar{W}_t^v(x, w) = \max_{y, p: y \geq x, \bar{p} \geq p \geq \underline{p}} E_{\tilde{\alpha}, \tilde{\beta}}[\bar{W}_{t+1}^v(x_+, w_+)],$$

where x_+ and w_+ are defined according to Eq. (14). Thus for the fixed v value, we have

$$G_x(v) = \bar{W}_1^v(x, 0).$$

Let $w' := w - v$ and $W_t^v(x, w') := \eta(\bar{W}_t^v(x, w' + v) - v)$, we obtain

$$W_{T+1}^v(x, w') = \min(w', 0)$$

and

$$W_t^v(x, w') = \max_{y, p: y \geq x, \bar{p} \geq p \geq \underline{p}} E_{\tilde{\alpha}, \tilde{\beta}}[W_{t+1}^v(x_+, w'_+)].$$

Notice that W_t^v does not depend on v. Thus W_t^v can be expressed as W_t and we have

$$\max_v G_x(v) = \max_v \Big(v + \frac{1}{\eta} W_1(x, -v)\Big),$$

which gives us the formulation in Section 4.1. Notice that the main advantage of this formulation, comparing to the DP algorithm on $\bar{W}_t^v$, is that we only need to solve the DP once since the external variable v is not involved in the DP recursion.

References

[1] Agrawal, V., S. Seshadri. 2000. Impact of Uncertainty and Risk Aversion on Price and Order Quantity in the Newsvendor Problem. *Manufacturing Service Operations Management*, **2**(4), pp. 410-423.

[2] Arrow, K. J. 1971. Essays in the theory of *Risk-Bearing*, Chicago: Markham.

[3] Bertsimas, D., G. J. Lauprete, A. Samarov. 2003. Shortfall as a risk measure: properties, optimization and applications. *Journal of Economics Dynamics and Control.*

[4] Bertsimas, D., A. Thiele. 2003. A Robust Optimization Approach to Supply Chain Management. Working Paper, MIT.

[5] Bourakiz, M., M. J. Sobel. 1992. Inventory Control with an Exponential Utility Criterion. *Operations Research*, **40**, pp. 603-608.

[6] Chen, F., A. Federgruen. 2000. Mean-Variance Analysis of Basic Inventory Models. Working paper, Columbia University.

[7] Chen, X. 2003. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost. Ph.D. thesis, Massechusetts Institute of Technology.

[8] Chen, X., D. Simchi-Levi. 2002. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case. To appear in *Operations Research.*

[9] Chen, X., D. Simchi-Levi. 2002. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Infinite Horizon Case. To appear in *Mathematics of Operations Research.*

[10] Chen, X., D. Simchi-Levi. 2003. A New Approach for the Stochastic Cash Balance Problem with Fixed Costs. Working paper, MIT.

[11] Dowd, K. 1998. *Beyond Value at Risk.* Weley, New York.

[12] Duffie, D., J. Pan. 1997. An overview of value at risk. Journal of Derivatives, 7-49.

[13] Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The Risk-Averse (and Prudent) Newsboy. *Management Science*, **41**(5), pp. 786-794.

[14] Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Operations Research*, **47**, No. 3, pp. 454-475.

[15] Huang, C. F., Litzenberger, R. H. 1998. *Foundations for Financial Economics.* Prentice Hall, Englewood Cliffs, NJ.

[16] Jorion, P. 1997. *Value at Risk.* McGraw-Hill, New York.

[17] Lau, H. S. 1980. The Newsboy Problem under Alternative Optimization Objectives. *Journal of the Operations Research Society*, **31**(1), pp. 525-535.

[18] Levy, H. 1992. Stochastic dominance and expected utility; survey and analysis. *Management Science*, **38** (4), 555-593.

[19] Levy, H., Y. Kroll. 1978. Ordering uncertain options with borrowing and lending. *Journal of Finance*, **31** (2).

[20] Markowitz, H. M. 1952. Portfolio Selection. *Journal of Finance*, **7**(1), 77-91.

[21] Martínez-de-Albéniz V., D. Simchi-Levi. 2003. Mean-Variance Trde-offs in Supply Contracts. Working paper, MIT.

[22] Mas-Collel, A., M. Whinston, J. Green. 1995. *Microeconomic Theory.* Oxford UP.

[23] Pratt, J. W. 1964. Risk Aversion in the Small and in the large. *Econometrica*, **32**, 122-136

[24] Rockafellar, R. T., S. Uryasev. 2000. Optimization of Conditional Value-at-Risk, *Journal of Risk*, **2**, 21-42.

[25] Rockafellar, R. T., S. Uryasev. 2002. Conditional Value-at-Risk for general loss distributions, *Journal of Banking and Finance*, **26**, 1443-1471.

[26] Scarf, H. 1960. The optimality of (s, S) policies for the dynamic inventory problem. *Proceedings of the 1st Stanford Symposium on Mathematical Methods in the Social Sciences*, Stanford University Press, Stanford, CA.

[27] Schweitzer, M., G. Chachon. 2000. Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence. *Management Science*, **46**(3), 404-420.