

Risk Aversion in Inventory Management*

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

October 3, 2004

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 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 functions is almost identical to the structure of the optimal risk-neutral inventory (and pricing) policies. These structural results are extended to models in which the decision maker has access to a (partially) complete financial market and can hedge its operational risk through trading financial securities. 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 managers 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, experimental evidence suggests that for some products, the so-called *high-profit* products, the decision makers are risk averse; see Schweitzer and Cachon [30] for more details. Unfortunately, traditional inventory control models fail short of meeting the needs of risk-averse planners. For instance, traditional inventory

*Research supported in part by the Singapore-MIT Alliance, ONR Contracts N00014-95-1-0232 and N00014-01-1-0146, and by NSF Contracts DMI-9732795, DMI-0085683 and DMI-0245352. The authors gratefully acknowledge many valuable suggestions from James E. Smith, the associate editor and two anonymous referees which led to significant improvement of the paper from its original version.

[†]Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign. Email: xinchen@uiuc.edu

[‡]NUS Business School, National University of Singapore. Email: dscsim@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

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 [21] analyzes the classical newsvendor model under two different objective functions. 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 [15] 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 [25] 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 [4] is closely related to ours. In this paper the authors 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 [2] 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.

Recently, we have seen a growing interest in hedging operational risk using financial instruments. As far as we know, all of this literature focus on single period (newsvendor) models with demand distribution that is correlated with the return of the financial market. This can be traced back to Anvari [1], which uses capital asset pricing model (CAPM) to study a newsvendor facing normal demand distribution. Chung [13] provides an alternative derivation for the result. More recently, Gaur and Seshadri [17] investigate the impact of financial hedging on the operations decision, and Caldentey and Haugh [5] show that different information assumptions lead to different types of solution techniques.

In this paper, we propose a general framework to incorporate risk aversion into multi-period inventory (and pricing) models. Specifically, 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. Following Smith [32], we take the standard economics perspective in which the decision maker maximizes the total expected utility from *consumption* in each time period. In Section 2 we discuss in more details the theory of expected utility employed in a multi-period decision making framework. We extend our framework (Section 4) by incorporating a *complete* or *partially complete* financial market so that the decision maker can hedge operational risk through financial securities.

Observe that if the utility functions are linear and increasing, the decision maker is risk-neutral and these problems are reduced to the classical finite horizon stochastic inventory problem and the

	<i>Price Not a Decision</i>		<i>Price is a Decision</i>	
	(Capacity) $k = 0$	$k > 0$	(Capacity) $k = 0$	$k > 0$
risk-neutral Model	(modified) base stock*	$(s, S)^*$	(modified) base stock list price*	$(s, S, A, p)^*$
Exponential Utility	(modified) base stock [†]	(s, S)	(modified) base stock	(s, S, A, p)
Increasing & Concave Utility	wealth dependent (modified) base stock	?	wealth dependent (modified) base stock	?
Partially Complete Financial Market Exponential Util.	state dependent (modified) base stock	state dependent (s, S)	state dependent (modified) base stock	state dependent (s, S, A, p)

Table 1: Summary of Previous Results and New Contributions. (*indicates existing results in the literature. †indicates a similar existing result based on a special case of our model.)

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 [29] 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 $s_t = S_t$ is the base stock level. This policy is optimal when $k = 0$. In addition, if there is a capacity constraint on the ordering quantity (expressed as “(Capacity)” in the table), then the modified base stock policy is optimal (expressed as “(modified)” in the table). That is, when the inventory level is below the base stock level, order enough to raise the inventory level to the base stock level if possible or order an amount equal to the capacity; otherwise, no order is placed.

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 [9]. 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. When A_t is empty for all t , we refer to such a policy as the (s, S, p) policy. 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. Again, when there is an ordering capacity constraint, a modified base stock inventory policy is optimal. See Federgruen and Heching [16] and Chen and Simchi-Levi [9].

The table suggests that when risk is measured using additive exponential utility functions, 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 $k = 0$. Notice that the optimal policy characterized by Bouakiz and Sobel [4] has the same structure as the optimal policy in our model. Finally, when $k = 0$ and there

is an ordering capacity constraint, a (modified) base stock policy is optimal.

When the risk averse inventory manager has access to a partially complete financial market to hedge its operations costs and part of the demand risks, the last row of Table 1 indicates that the structure of the optimal inventory control policy when the decision maker’s utility function is exponential. If the financial market is *complete* instead of partially complete, our model reduces to the risk neutral case and hence we have the same structural results as the risk neutral model with respect to the market risk neutral probability. We will explain the meaning of “state dependent” when we present the model in Section 4.

We complement the theoretical results with a numerical study demonstrating that this framework can identify the tradeoff between expected profit and the risk of under-performing. Finally, we empirically illustrate the effect of risk aversion on the inventory policies.

The paper is organized as follows. In Section 2, we review classical expected utility approaches in risk-averse valuation. In Section 3, we propose a model to incorporate risk aversion in multi-period inventory (and pricing) setting, and focus on characterizing the optimal inventory policy for a risk-averse decision maker. We then generalize the results in Section 4 by considering the financial hedging option. Section 5 presents the computational results illustrating the effects of different risk-averse multi-period inventory models on profit distribution and inventory control policies. Finally, Section 7 provides some concluding remarks.

We complete this section with a brief statement on notations. Specifically, a variable with tilde over it such as \tilde{d} , denotes a random variable.

2 Utility theory for risk-averse valuations

Modeling risk sensitive decision making is one of the fundamental problems in economics. A basic theoretical framework for risk sensitive decision making is the so called expected utility theory (see, e.g., [26], Chapter 6).

Assume a decision maker has to make a decision in a single period problem before uncertainty is resolved. According to the expected utility theory, the decision maker’s objective is to maximize the expectation of some appropriately chosen utility function of the decision maker’s payoff. Such a modeling framework for risk sensitive decision making is established mathematically based on an axiomatic argument. That is, based on a certain set of axioms regarding the decision maker’s preference relation, one can show the existence of such a utility function and that the decision maker’s choice criterion is the expected utility. See, for example, Heyman and Sobel [19] Chapter 2-4, and Fishburn [18].

For multi-period problems, one approach of modeling risk aversion that seems natural is to maximize the expected utility of the present value of the income cash flow. In calculating the present value, one may take the interest rate for risk free borrowing and lending as the discount factor, reflecting the fact that the decision maker could borrow and lend over time and convert any deterministic cash flow into its present value. Models based on this approach are referred to as the *present value models*. Note that the present value models have been employed by Bouakiz and Sobel [4] and Chen, Sim, Simchi-Levi and Sun [8] to analyze the multi-period inventory replenishment problems of a risk averse inventory manager.

However, in the economics literature it has long been known that applying expected utility methods directly to income cash flows causes the so called “temporal risk problem” — it does not capture the decision maker’s sensitivity to the time at which uncertainties are resolved, see, e.g.,

a summary description of this problem in Smith [32]. One way to overcome the “temporal risk problem” is to explicitly model the utility over a flow of consumption, allowing the decision maker to borrow and lend to “smooth” the income flow as the uncertainties unfold over time.

Therefore, an alternative modeling approach for the multi-period inventory control problem is to directly model consumption, saving and borrowing decisions as well as inventory replenishment and pricing decisions.

Assume that the decision maker has access to a financial market for borrowing and lending with a risk free saving and borrowing interest rate r_f . At the beginning of period t , assume that the decision maker has initial wealth w and chooses an operational policy to affect her income flow. While at the end of period t , that is, after the uncertainty of this period has been resolved, the decision maker observes her current wealth level $w + \bar{P}_t$ and decides her consumption level f_t for this period, where \bar{P}_t is the income generated at period t . The remaining wealth, $w + \bar{P}_t - f_t$, is then saved (or borrowed, if negative) for the next period. Thus, the next period’s initial wealth is

$$z = (1 + r_f)(w + \bar{P}_t - f_t) ,$$

where $v^+ = \max(x, 0)$ and $v^- = \min(x, 0)$. Equivalently, we can model z as the decision variable and thus,

$$f_t = w - \frac{z}{1 + r_f} + \bar{P}_t .$$

The decision maker’s objective is to maximize her expected utility of the consumption flow,

$$E[U(f_1, \dots, f_T)] ,$$

over the planning horizon $1, \dots, T$. We call such an approach the *consumption model*. Smith [32] provides an excellent comparison between the consumption model and the present value model.

Similar to single period problems, axiomatic approaches were also employed to derive certain types of utility functions for multi-period problems (see, e.g., Sobel [34] and [23], Chapter 9). In particular, the so called “utility independence axiom” implies the following additive utility functions, that is, the utility of the consumption flow is the summation of the utility from the consumption from each period

$$U(f_1, f_2, \dots, f_T) = \sum_{t=1}^T u_t(f_t) ,$$

where u_t is increasing and concave. As a special case of the additive utility functions, the additive exponential utility functions are also commonly used in economics ([26]). In this case, the utility function has the form $u_t(f_t) = -a_t e^{-f_t/\rho_t}$ for some parameters $a_t > 0$ and $\rho_t > 0$. Howard [20] indicates that exponential utility functions are widely applied in decision analysis practice. Kirkwood [24] shows that in most cases, an appropriately chosen exponential utility function is a very good approximation for general utility functions.

In the next section we characterize the structure of the inventory policies according to the consumption model. Interestingly, when the additive exponential utility functions are used, the consumption model reduces to the present value model, which implies that the structure of the inventory policies of the present value model is the same as that of the consumption model.

At this point it is worth mentioning that Savage [28] unified the von Neumann and Morgenstern’s theory of expected utility and de Finetti’s theory of subjective probability and established the *subjective expected utility theory*. Without assuming probability distributions and utility functions, the

Savage theory starts from a set of assumptions on the decision maker's preferences and shows the existence of a (subjective) probability distribution depending on the decision maker's belief on the future state of the world as well as a utility function. The decision maker's objective is to maximize the expected utility, with the expectation taken according to the subjective probability distribution. In Section 4, in order to introduce the framework of risk averse inventory management with financial hedging opportunities, we explicitly consider the decision maker's subjective probability and distinguish it from the so called *risk neutral probability* reflected by a (partially) complete financial market with no arbitrage opportunity. A similar approach has been employed by Smith and Nau [33] and Gaur and Seshadri [17].

3 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 , $t = 1, 2, \dots$, 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 $t = 1, 2, \dots$, the demand function satisfies

$$\tilde{d}_t = D_t(p_t, \tilde{\epsilon}_t) := \tilde{\beta}_t - \tilde{\alpha}_t p_t, \quad (1)$$

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 , $t = 1, 2, \dots$, 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 $\lim_{|x| \rightarrow \infty} h_t(x) = \infty$.

At the beginning of period t , the inventory manager decides the order up to level y_t and the price p_t . After observing the demand she then makes consumption decision f_t . Thus, given the initial inventory level x_t , the order up to level y_t and the realization of the uncertainty ϵ_t , the income at period t is

$$\bar{P}_t(x_t, y_t, p; \tilde{\epsilon}_t) = -k\delta(y_t - x_t) - c_t(y_t - x_t) + p_t D_t(p_t, \tilde{\epsilon}_t) - h_t(y_t - D_t(p_t, \tilde{\epsilon}_t)).$$

Moreover, as discussed in the previous section, the consumption decision at period t is equivalent to deciding on the initial wealth level of period $t + 1$. Let w_t be the initial wealth level at period t . Then

$$f_t = w_t - \frac{w_{t+1}}{1 + r_f} + \bar{P}_t(x_t, y_t, p_t; \tilde{\epsilon}_t).$$

Finally, at the last period T , we assume the inventory manager consumes everything, which corresponds to $w_{T+1} = 0$.

According to the consumption model, the inventory manager's decision problem is to find the order-up-to levels y_t , the selling price p_t and decide the initial wealth level w_t (correspondingly, the consumption level) to the following optimization problem.

$$\begin{aligned} \max \quad & E[U(f_1, f_2, \dots, f_T)] \\ \text{s.t.} \quad & y_t \geq x_t, \\ & x_{t+1} = y_t - D_t(p_t, \tilde{\epsilon}_t), \\ & f_t = w_t - \frac{w_{t+1}}{1+r_f} + \bar{P}_t(x_t, y_t, p_t, \epsilon_t), \\ & w_{T+1} = 0. \end{aligned} \tag{2}$$

When the utility function $U(f_1, f_2, \dots, f_T)$ takes the following form,

$$U(f_1, f_2, \dots, f_T) = \sum_{t=1}^T \frac{f_t}{(1 + r_f)^{t-1}},$$

the consumption model reduces to the traditional risk-neutral inventory (and pricing) problem analyzed by Chen and Simchi-Levi [9]. In this case, 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 (See Chen and Simchi-Levi [9] for more details).

$$V_t(x) = c_t x + \max_{y \geq x, \underline{p}_t \leq p \leq \bar{p}_t} -k\delta(y - x) + g_t(y, p), \tag{3}$$

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

$$g_t(y, p) = E\{pD_t(p, \tilde{\epsilon}) - c_t y - h_t(y - D_t(p, \tilde{\epsilon})) + \frac{1}{1 + r_f} V_{t+1}(y - D_t(p, \tilde{\epsilon}))\}. \tag{4}$$

The following theorem presents known results for the traditional risk-neutral models.

Theorem 3.1 (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 $\max_{\underline{p}_t \leq p \leq \bar{p}_t} g_t(y, p)$ 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 [29] using the concept of k -convexity; part (b) and part (c) are proved in Chen and Simchi-Levi [9] 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 [9] hold true under more general demand functions than those in Assumption 1.

In the following subsections, we analyze the consumption model based on the additive utility functions and its special case the additive exponential utility functions.

3.1 Additive utility model

In this subsection, we focus on the additive utility functions. According to the sequence of events as described before, the optimization model (2) can be solved by the following dynamic programming recursion.

$$V_t(x, w) = \max_{y \geq x, p_t \leq p \leq \bar{p}_t} E_{\tilde{\epsilon}} \left[\bar{W}_t(x, w, y, p; \tilde{\epsilon}_t) \right], \quad (5)$$

in which

$$\bar{W}_t(x, w, y, p; \tilde{\epsilon}) = \max_z \left\{ u_t \left(w - \frac{z}{1 + r_f} + \bar{P}(x, y, p; \tilde{\epsilon}_t) \right) + V_{t+1}(y - D_t(p, \tilde{\epsilon}_t), z) \right\} \quad (6)$$

with boundary condition $V_{T+1}(x, w) = 0$.

Notice that unlike the traditional risk-neutral inventory models, where the state variable in the above dynamic programming recursion is the current inventory level, here we augment the state space by introducing a new state variable, namely, the wealth level w .

Instead of working with the dynamic program (5)-(6), we find that it is more convenient to work with an equivalent formulation. Let

$$U_t(x, w) = V_t(x, w - c_t x),$$

and

$$P_t(y, p; \tilde{\epsilon}_t) = \left(\frac{c_{t+1}}{1 + r_f} - c_t \right) y + \left(p - \frac{c_{t+1}}{1 + r_f} \right) D_t(p, \tilde{\epsilon}_t) - h_t(y - D_t(p, \tilde{\epsilon}_t)).$$

The dynamic program (5)-(6) becomes

$$U_t(x, w) = \max_{y \geq x, p_t \leq p \leq \bar{p}_t} E \left[W_t(x, w, y, p; \tilde{\epsilon}_t) \right], \quad (7)$$

in which

$$W_t(x, w, y, p; \tilde{\epsilon}_t) = \max_z \left\{ u_t \left(w - \frac{z}{1 + r_f} - k\delta(y - x) + P_t(y, p; \tilde{\epsilon}_t) \right) + U_{t+1}(y - D_t(p, \tilde{\epsilon}_t), z) \right\} \quad (8)$$

with boundary condition $U_{T+1}(x, w) = 0$.

Theorem 3.2 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 (w) dependent base stock inventory policy is optimal.

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 $\tilde{\epsilon}_t$, P_t is jointly concave in (y, p) . Thus,

$$W_t(w, y, p; \tilde{\epsilon}_t) = \max_z \left\{ u_t \left(w - \frac{z}{1+r_f} + P_t(y, p; \tilde{\epsilon}_t) \right) + U_{t+1}(y - D_t(p, \tilde{\epsilon}_t), z) \right\}$$

is jointly concave in (w, y, p) , which further implies that $E[W_t(w, y, p; \tilde{\epsilon}_t)]$ is jointly concave in (w, 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

$$\max_y \left\{ \max_{\bar{p}_t \geq p \geq \underline{p}_t} E[W_t(w, y, p; \tilde{\epsilon}_t)] \right\}.$$

Since $E[\bar{W}_t(w, y, p; \tilde{\epsilon}_t)]$ 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. ■

Theorem 3.2 can be extended to incorporate capacity constraints on the order quantities. In this case, it is straightforward to see that the proof of Theorem 3.2 goes through. The only difference is that in this case, a w -dependent modified base stock policy is optimal. In such a policy, when the initial inventory level is no more than $y^*(w)$, order up to $y^*(w)$ if possible; otherwise order up to the ordering capacity. On the other hand, no order is placed when the initial level is above $y^*(w)$.

Recall that in the case of a risk-neutral decision maker, a base stock list price policy is optimal. Theorem 3.2 thus implies that the optimal inventory policy for the expected additive utility risk-averse model is quite different. Indeed, in the risk-averse case, the base stock level depends on the wealth, measured by the position of the risk free financial security. Moreover, it is not clear in this case whether a list price policy is optimal or the banking decisions has any nice structure.

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

3.2 Exponential utility functions

We now focus on exponential utility functions of the form $u_t(f) = -a_t e^{-f/\rho_t}$ with parameters $a_t, \rho_t > 0$. ρ_t is the risk tolerance factor while a_t 's reflect the decision maker's attitude towards the utility obtained from different period.

The beauty of exponential utility functions is that we are able to separately make the inventory decisions without considering the banking/consumption decisions. This is discovered by Smith [32] in the decision tree framework. The following theorem states this result in the dynamic programming language. For completeness, a proof is presented in Appendix A.

First, we introduce the notation of "effective risk tolerance" per period defined as

$$R_t = \sum_{\tau=t}^T \frac{\rho_\tau}{(1+r_f)^{\tau-t}}. \quad (9)$$

Theorem 3.3 Assume that $u_t(f) = -a_t e^{-f/\rho_t}$. The inventory decisions in the risk-averse inventory control model Eq. (5)-(6) can be calculated through the following dynamic programming recursion

$$G_t(x) = \max_{y \geq x, \bar{p}_t \geq p \geq \underline{p}_t} -k\delta(y-x) - L_t(y, p), \quad (10)$$

where

$$L_t(y, p) = R_t \ln E \left[\exp \left\{ -\frac{1}{R_t} \left[P_t(y, p; \tilde{\epsilon}_t) + \frac{1}{1+r_f} G_{t+1}(y - D_t(p, \tilde{\epsilon}_t)) \right] \right\} \right], \quad (11)$$

and $G_{T+1}(x) = 0$.

The optimal consumption decision at each period $t = 1, \dots, T-1$ is

$$f_t^*(x, y, p, \tilde{d}) = \frac{\rho_t}{R_t} \left[w + (-k\delta(y-x) + P_t(y, p; \tilde{\epsilon}_t)) + \frac{1}{1+r_f} G_{t+1}(y - \tilde{d}) \right] + C_t,$$

in which C_t is a constant that does not depend on (x, y, p, \tilde{d}) .

The theorem thus implies that when additive exponential utility functions are used (i) the optimal inventory policy is independent of wealth level; (ii) the optimal inventory replenishment and pricing decisions can be obtained regardless of the banking/consumption decisions; (iii) the optimal consumption decision is a simple linear function of the current wealth level; and (iv) the model parameter a_t does not affect the inventory replenishment and pricing decisions. Thus, incorporating additive exponential utility function significantly simplifies the problem. We also point it out that $\ln E \exp(\cdot)$ is the so called certainty equivalent operator in decision analysis literature.

This theorem, together with Theorem 3.2, 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. As we show at the end of this subsection, the present value model is a special case of the consumption model. Thus, our results for the case when price is not a decision variable directly imply the results obtained by Bouakiz and Sobel [4] using a more complicated argument. If in addition, there is a capacity constraint on ordering, one can show that a wealth independent modified base stock policy is optimal. 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 proposition.

Proposition 1 If a function f is convex, k -convex or symmetric k -convex, then the 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

$$Z(x_\lambda, w_\lambda) \leq (1 - \lambda)Z(x_0, w_0) + \lambda Z(x_1, w_1), \quad (12)$$

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

Dividing both sides of (12) 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 the additive exponential utility function.

Theorem 3.4 (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 Proposition 1, $G_t(y, p)$ is k -concave (or symmetric k -concave). The remaining parts follow directly from Lemma 1 and Proposition 2 for k -concavity or Lemma 2 and Proposition 3 for symmetric k -concavity. Lemma 1, Proposition 2, Lemma 2 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 3.1 part (a). Theorem 3.4 implies that this is also true under the exponential utility measure. Similarly, for the more general inventory and pricing problem, Theorem 3.1 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 3.1 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.

Now we show that if the additive exponential utility function is used, the present value model is a special case the consumption model and thus the same structure results in Theorem 3.4 holds for the present value model. Indeed, if $\rho_t = 0$ for $t = 1, \dots, T - 1$, the consumption model (5)-(6) reduces to the present value model. The intuition is clear. In fact, $\rho_t = 0$ implies that the decision maker is “extremely risk-averse” and thus the consumption has to be non-negative. On the other

hand, since for any non-negative consumption level in period $t = 1, \dots, T - 1$, the utility is always 0, the decision maker is better off by shifting the consumption to another period. Eventually, the optimal consumption decisions are to save all income flow to the last period. Thus, in this case, the consumption model reduces to the present value model.

4 Multi-period Inventory Models with Financial Hedging Opportunities

In this section, we extend our previous framework by assuming that the decision maker has opportunities of hedging operational risk through financial securities.

Similar to the previous section, we consider a risk-averse inventory manager that has to make replenishment (and pricing) decisions over a finite time horizon with T periods. Following Song and Zipkin [35], we assume that at each time period, the business environment could be in one of a number of possible levels. Inventory model parameters and the sufficient statistics of the demand distribution depends on the history of the evolution of the business environment. Formally, let finite set Θ_t represent the set of period t business environments. And we denote bold case $\Theta_t = \prod_{\tau=1}^t \Theta_\tau$ to be the set of trajectories of levels from period 1 to t (also referred to as *state of the world* in the sequel). Its components uniquely determine the cost structure and the demand distribution of the inventory model and the price of the securities in the financial market. That is, at each time period t parameters k , c_t , h_t , \underline{p}_t and \bar{p}_t are all functions of θ_t (we express them as $k_t^{\theta_t}$, $c_t^{\theta_t}$, $h_t^{\theta_t}$, $\underline{p}_t^{\theta_t}$ and $\bar{p}_t^{\theta_t}$), and the (decision maker's subjective probability) distributions of $\tilde{\alpha}$ and $\tilde{\beta}$ are also θ_t -dependent. As before, we denote $\tilde{\epsilon}^{\theta_t} - t = (\alpha_t^{\theta_t}, \beta_t^{\theta_t})$. Notice that in this section we explicitly distinguish subjective probabilities from the probability distribution that can be inferred from the financial market, known as the “risk-neutral probability”, which will be introduced next.

Thus the decision maker's income at period t is

$$\bar{P}_t(x_t, y_t, p_t, \tilde{\epsilon}; \theta_t) = -k_t^{\theta_t} \delta(y_t - x_t) - c_t^{\theta_t} (y_t - x_t) + p_t D_t(p_t, \tilde{\epsilon}^{\theta_t}) - h_t^{\theta_t} (y_t - D_t(p_t, \tilde{\epsilon}^{\theta_t})) . \quad (13)$$

We now describe the financial market and the so called *partially complete* assumption.

We assume that the risk free interest rate for borrowing or lending “cash” is r_f . That is, one dollar cash in period t becomes $1 + r_f$ dollars in period $t + 1$ under all states of the world evolution. We assume that there are N financial securities in the financial market. To simplify notation and analysis, assume that these securities do not pay dividends during the time horizon $1, \dots, T$. We denote the prices of the securities as a matrix Q such that its component Q_{it} denotes the price of security i at time t (measured by period t dollar). Following Smith and Nau [33], we assume the levels of model parameters are the so called *market risk* – it can be fully hedged by the financial market. On the other hand, given the state of the world, the demand uncertainty is the so called *private risk* that cannot be hedged by the financial market. Formally, this implies the following conditions

- Assumption 2** (1) Q_{it} only depends on the state of the world $\theta_t \in \Theta_t$. That is, for any trajectory $\theta_T = \{\vartheta_1, \dots, \vartheta_T\}$ and its sub-trajectories $\theta_t = \{\vartheta_1, \dots, \vartheta_t\}$, we can uniquely define the price sequence of financial security i as the (row) vector $Q_{i \cdot}(\theta_T)$.
- (2) Any cash flow determined by the state of the world can be replicated by trading the financial securities. That is, for any given period t , the vector $\{\mathcal{P}_t(\theta_t)\}_{\forall \theta_t \in \Theta_t}$ is a linear combination of

the vectors $\mathbf{1}$ and $\{Q_{1t}(\theta_t)\}_{\forall \theta_t \in \Theta_t}, \dots, \{Q_{Nt}(\theta_t)\}_{\forall \theta_t \in \Theta_t}$. Here $\mathbf{1}$ is the all one vector and $\mathcal{P}_t(\cdot)$ is any mapping from Θ_t to real number, with $\bar{P}_t(x_t, y_t, p_t, \tilde{c}_t^{\theta_t}; \cdot)$ in Eq. (13) being a special case.

- (3) Disclosed demand information in each time period is not correlated to any future evolution of the state of the world. That is, given θ_t , the decision maker believes that $d_t^{\theta_t}$ and θ_{t+1} are independent.

We also assume that

Assumption 3 The financial market is arbitrage-free. That is, there does not exist a sequence of $N + 1$ dimensional vectors $\{(w_t, \mathbf{w}_t)\}_{t=1, \dots, T}$ such that

$$\sum_{t=1}^T \frac{w_t}{(1+r_f)^{t-1}} + \sum_{t=1}^T \mathbf{w}_t^\top Q_{\cdot 1} \leq 0$$

and for all t and θ_t ,

$$w_t + \mathbf{w}_t^\top Q_{\cdot t}(\theta_t) \geq 0$$

with $w_t + \mathbf{w}_t^\top Q_{\cdot t}(\theta_t) > 0$ for some t .

Assumption 2, parts (1) and (2), also implies an equivalent “dual” characterization of the no arbitrage assumption – a security market is *arbitrage free* if and only if there exists a strictly positive probability distribution π (commonly referred to as the *risk-neutral probability*) on the states of the world Θ such that for all t ,

$$Q_{i1} = \frac{1}{(1+r_f)^{t-1}} \sum_{\theta_t \in \Theta_t} \pi(\theta_t) Q_{it}(\theta_t),$$

in which $\pi(\theta_t)$ is the risk neutral probability of observing the trajectory θ_t , that is, $\pi(\theta_t) = \sum_{\theta_T: \theta_t \in \theta_T} \pi(\theta_T)$, where we use $\theta_T : \theta_t \in \theta_T$ to express all the trajectories $\theta_T \in \Theta_T$ that contains θ_t as the first t period sub-trajectory.

In the sequel we use $E_\pi[\cdot|\theta_t]$ to denote the conditional expectation taken with respect to the risk neutral probability distribution π while $E_{\theta_{t+1}}[\cdot|\theta_t]$ is used to express the expectation taken with respect to the decision maker’s subjective probability. When we take expectation on the subjective demand distribution, we use notation $E_{\tilde{c}_t^{\theta_t}}[\cdot]$.

We use N dimensional vector \mathbf{w}_t to express the inventory manager’s financial market position at time period t . That is, at time period t , her holding of security i is $\mathbf{w}_t(i)$. We use scalar w_t to represent the amount of cash in the bank at period t . At each time period t , the decision maker observes the current state of the world θ_t , the current inventory level x_t and the current financial market position \mathbf{w}_t , then make the inventory and pricing decisions y_t and p_t . After observing the realized demand (and thus the income cash flow P_t), she makes the decision on the next period market position \mathbf{w}_{t+1} by trading at the market price $Q_{\cdot t}$. With the amount f_t consumed for utility at period t , the period $t + 1$ cash amount becomes

$$w_{t+1} = (1+r_f) \left(w_t + \bar{P}_t(x, y, p, \tilde{c}_t^{\theta_t}; \theta_t) + (\mathbf{w}_t - \mathbf{w}_{t+1})^\top Q_{\cdot t} - f_t \right).$$

Equivalently, we have

$$f_t(w_t, w_{t+1}, \mathbf{w}_t, \mathbf{w}_{t+1}, x, y, p, \tilde{c}_t^{\theta_t}; \theta_t) = (\mathbf{w}_t - \mathbf{w}_{t+1})^\top Q_{\cdot t} + \bar{P}_t(x, y, p, \tilde{c}_t^{\theta_t}; \theta_t) + w_t - \frac{w_{t+1}}{1+r_f}.$$

The objective of the inventory manager is to find an ordering (and pricing) policy as well as a hedging strategy so as to maximize her expected utility over consumptions. This maximization problem can be expressed by the following dynamic programming recursion

$$V_t(x, w, \mathbf{w}, \theta_t) = \max_{y, p: y \geq x, \underline{p}_t^{\theta_t} \leq p \leq \bar{p}_t^{\theta_t}} E_{\tilde{\epsilon}_t^{\theta_t}} \left[\bar{W}_t(x, w, \mathbf{w}, y, p; \tilde{\epsilon}_t^{\theta_t}, \theta_t) \right], \quad (14)$$

in which

$$\begin{aligned} \bar{W}_t(x, w, \mathbf{w}, y, p; \tilde{\epsilon}_t^{\theta_t}, \theta_t) = & \max_{z, \mathbf{z}} \left\{ u_t \left(f_t(w, z, \mathbf{w}, \mathbf{z}, x, y, p, \tilde{\epsilon}_t^{\theta_t}; \theta_t) \right) \right. \\ & \left. + E_{\theta_{t+1}} \left[V_{t+1}(y - D_t(p, \tilde{\epsilon}_t^{\theta_t}), z, \mathbf{z}, \theta_{t+1}) | \theta_t \right] \right\} \end{aligned} \quad (15)$$

with boundary condition $V_{T+1}(x, w, \mathbf{w}, \theta) = 0$.

Notice that all the expectations taken in the above dynamic programming model are with respect to the decision maker's subjective probabilities.

A special case of the partially complete market assumption is obtained when $\tilde{\epsilon}^{\theta_t}$ is deterministic for any given θ_t . This corresponds to the *complete market* assumption. Following Smith and Nau [33] we know that a risk averse inventory manager with additive concave utility function can fully hedge the risk in a complete market while locking in a profit equal to the expected (with respect to the risk neutral probability) profit. Thus, in this case, the inventory control problem reduces to a risk neutral problem.

On the other hand, under the partially complete market assumption, the following theorem holds for a decision maker with the additive exponential (subjective) expected utility maximization criterion. This theorem can be obtained directly from Section 5 of Smith and Nau [33] and the previous section of this paper. For simplicity, let

$$P_t(y, p, \tilde{\epsilon}_t; \theta_t) = \bar{P}_t(x, y, p, \tilde{\epsilon}_t; \theta_t) + k_t^{\theta_t} \delta(y - x) - c_t^{\theta_t} x.$$

Theorem 4.1 *The inventory and pricing decisions in the risk-averse inventory model with financial hedging Eq. (14)-(15) can be calculated through the following dynamic programming recursion*

$$G_t(x, \theta_t) = c_t^{\theta_t} x + \max_{y, p: y \geq x, \underline{p}_t^{\theta_t} \leq p \leq \bar{p}_t^{\theta_t}} -k_t^{\theta_t} \delta(y - x) + L_t(y, p, \theta_t), \quad (16)$$

in which

$$L_t(y, p, \theta_t) = -R_t \ln E_{\tilde{\epsilon}_t^{\theta_t}} \left[\exp \left\{ - \left(P_t(y, p, \tilde{\epsilon}_t; \theta_t) + \frac{1}{1 + r_f} E_{\pi} \left[G_{t+1}(y - D_t(p, \tilde{\epsilon}_t^{\theta_t}), \theta_{t+1}) | \theta_t \right] \right) / R_t \right\} \right]$$

and with boundary condition $G_{T+1}(x) = 0$.

If in addition, $k_t^{\theta_t} \geq E_{\pi} [k_{t+1}^{\theta_{t+1}} | \theta_t]$, then the following structural results for the optimal inventory (pricing) policies holds.

- (a) If price is not a decision variable (i.e., $\underline{p}_t^{\theta_t} = \bar{p}_t^{\theta_t}$ for each t), for each given θ_t , functions $G_t(x, \theta)$ and $L_t(y, p, \theta_t)$ are $k_t^{\theta_t}$ -concave and an θ_t dependent (s, S) inventory policy is optimal.
- (b) For the general case, $G_t(x, \theta_t)$ and $L_t(y, p, \theta_t)$ are symmetric $k_t^{\theta_t}$ -concave for any given θ_t and an θ_t dependent (s, S, A, p) policy is optimal.

The theorem thus implies that when additive exponential utility functions are used (i) the optimal inventory policy is independent of the financial market position; (ii) the optimal inventory replenishment and pricing decisions can be obtained regardless of the financial hedging decisions; (iii) the coefficient a_t in the utility function does not affect the inventory replenishment and pricing decisions; and (iv) unlike equation (15), the expectation operator $E_{\theta_{t+1}}[\cdot|\theta_t]$ does not appear in the above dynamic programming recursion, which implies that for the purpose of calculating the optimal inventory decisions, we do not need to know the decision maker's subjective probability on the state of the world evolution. However, in order to obtain the optimal expected utility, the model requires that the decision maker also implement an optimal strategy on the financial market. We refer readers to Smith and Nau [33] for the detailed description of such an optimal trading strategy.¹

It is appropriate to point out that the restriction on fixed costs in the theorem is similar to assumptions made in Sethi and Cheng [31] for a stochastic inventory model with input parameters driven by a Markov chain. Finally, when the fixed costs are all zeros and there are capacity constraints on the ordering quantities, our analysis shows that a state dependent modified base stock policy is optimal.

5 Computational Results

In this section, we present the results of an empirical study. We consider an additive exponential utility model in which $\rho_t = \rho$ for all $t = 1, \dots, T$. Assuming that the risk free interest rate, $r_f = 0$, the experimental model focuses on how the choice of parameter ρ can affect the entire inventory replenishment policies.

We experimented many different demand distributions and inventory scenarios and observed similar trends in profit profile and inventory policy changes under the influence of risk aversion. Hence, we highlight a typical experimental setup in which 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,

$$\tilde{d} = \min(\max(\lfloor 30\tilde{z} \rfloor + 10, 0), 150),$$

where $\tilde{z} \sim \mathcal{N}(0, 1)$, and $\lfloor y \rfloor$, the floor function, denotes the largest integer smaller or equal to y . Since the demand distribution is bounded and discrete, we can easily evaluate expectations within the dynamic programming recursion and compute the optimal policy exactly.

To evaluate the inventory policies derived, we analyze the profit distributions at the T period via Monte Carlo simulation on S independent trials. In each trial, we generate T independent demand samples (one for each period) and obtain the accumulated profit at the end of the T th period. Hence, in the policy evaluation stage, we require ST independent demands drawn from \tilde{d} . We can improve the resolutions of the policy evaluation by increasing the number of independent trials, S . Hence, the choice of S is limited by computation time, and in our experiment we choose $S = 10,000$.

¹We caution reader on the difference in the notation of this paper and the Smith and Nau [33]. In this paper we measure the prices of financial securities in the period t dollar, while Smith and Nau measure in period 1 dollar.

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

ρ	Mean Profit	Standard Deviation of Profit	% Negative Profit	Mean Loss
10	253	209	10.01	66.6
20	258	215	10.19	71.4
40	261	227	12.02	81.1
risk-neutral	266	234	12.26	89.8

Table 3: Summary of results for exponential utility model.

For each risk parameter $\rho \in \{10, 20, 40\}$, we construct the optimal risk-averse inventory policy. We also compare the profit profile of the risk-neutral replenishment policy. Table 3 shows the profit profile as the parameter ρ 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 ‘risk-neutral’ represents the risk-neutral case.

Observe that as the parameter ρ decreases, the sample mean profit and standard deviation of the profit profile decrease, indicating greater risk-averse decisions. At the same time, the percentage of negative profit and the mean loss decreases as well, i.e., decreases as the decision maker becomes more risk averse.

We now study numerically how the replenishment policies change as we vary the risk aversion level. That is, since the optimal policy is (s_t, S_t) , we analyze changes in the replenishment policy parameters as we vary the decision maker risk aversion level. Figure 1 depicts the parameters (s_t, S_t) over the first 9 periods. Generally, for any time period, the order-up-to level, S_t , decreases in response to greater risk aversion.

Interestingly, for this particular problem instance, the reorder level, s_t , increases as we increase the level of risk aversion. Of course, this is not true in general. As a matter of fact, if the fixed ordering cost, $k = 0$, we have $s_t = S_t$, and unless the policies are indifferent to risk aversion, we do not expect such phenomenon to hold. Indeed, it is not difficult to come up with examples (with different values of k) showing that s_t decreases in response to greater risk aversion. We point out that in most of our experiments, the order-up-to level S_t decreases, while the reorder points s_t are monotonic (both monotone increases and decreases are possible) in response to greater risk aversion. Unfortunately, while such monotonicity property is much desired, we have numerical examples that violate this property as we change the risk aversion level.

To test the sensitivity of the parameters of the optimal policy to changes in the level of risk aversion, we track the changes in the parameters of the optimal policy as we gradually increase the parameter ρ . It is interesting to observe that the policy changes are quite small. For instance, the




Figure 1: Plot of (s_t, S_t) against t .

optimal policy remains the same as we vary $\rho = 15, 16, \dots, 24$. Therefore, we conclude (numerically) that the optimal policy is relatively insensitive to small changes in the decision maker's level of risk aversion.

6 Effect 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. The numerical experiment of Bertsimas and Thiele [3] exemplifies the phenomenon that using the wrong demand distribution can substantially reduce the expected profit compared to the expected profit associated with the true optimal inventory policy. Evidently, inaccurate estimates of demand distributions yield inappropriate replenishment and pricing decisions, leading to poor inventory policies. Hence, given limited knowledge of demand distributions, there is an implicit “risk” of obtaining poor inventory policies. Our objective in this section is two-fold. First, we illustrate that models based on limited demand information may indeed lead to poor inventory policies. Second, we show that risk-averse models can alleviate to some extent the “risk” of obtaining poor inventory policies due to limited demand information.

The numerical experiment setup is similar to what we did in Section 5, namely we use additive exponential utility functions, same demand distribution across different periods, price is fixed and cost parameters are the same as in Table 2. However, we assume that the underline demand distribution is unknown to the decision maker. Instead, the decision maker has N observations of the demand data, d_1, \dots, d_N , or equivalently, $\epsilon_1, \dots, \epsilon_N$. These observations, representing historical data, have been generated independently from the underline demand distribution. Following machine learning terminology, we refer to this data as the *training data*. The small number N represents the situation

in which one has limited knowledge of the demand distributions. Since the demand distributions at every period are identical, the decision maker utilizes the same set of training data to obtain the replenishment policy at each period.

We now describe how the decision maker constructs the inventory policy when she has only limited demand information. We illustrate the approach through the dynamic programming recursion in (10)-(11). Since the decision maker has limited demand information, instead of solving the dynamic programming recursion in (10)-(11) in which the underline demand distribution is required, she constructs the policy that maximizes the estimated expected profit-to-go using the training data as follows

$$\hat{G}_t(x) = \max_{y \geq x} -k\delta(y-x) - \hat{L}_t(y, p_t),$$

where p_t denotes the fixed price of period t , and

$$\hat{L}_t(y, p) = R_t \ln \left\{ \frac{1}{N} \sum_{i=1}^N \left[\exp \left\{ -\frac{1}{R_t} \left[P_t(y, p; \epsilon_i) + \frac{1}{1+r_f} \hat{G}_{t+1}(y - D_t(p, \epsilon_i)) \right] \right\} \right] \right\},$$

i.e., we replace the expectations by their corresponding sample means. Similarly, risk-neutral inventory policies based on the limited demand information can be also constructed by following the same approach, i.e., replace the expectations in the dynamic programming recursion by the corresponding sample mean based on the observed data.

First, we compare two risk-neutral inventory policies, the risk-neutral inventory policy based on limited demand information (LD policy for short) with $N = 100$ and the risk-neutral inventory policy based on the underline full demand distribution (the true optimal inventory policy, FD policy for short). The first policy is derived using the approach described in the previous paragraph, while the second policy is derived using the underline demand distribution as in the previous section. For each policy, we can generate the sample profit via Monte Carlo with S samples as described in the previous section.

Using the LD policy, the profit mean and standard deviation are respectively 250 and 272, while employing the FD policy, the mean and standard deviation improve (that is, higher mean and lower standard deviation) to respectively, 266 and 234.

In addition to comparing the profit profiles using means and standard deviations, we will show that the sample profit profile generated by the risk-neutral LD policy is second order stochastically dominated by the profit profile generated by the risk-neutral FD policy. In other words, letting v_1, \dots, v_S and u_1, \dots, u_S be the profit independent samples derived using the risk-neutral LD policy and the risk-neutral FD policy respectively, we will show that

$$\hat{E}(f(\tilde{v})) = \frac{1}{S} \sum_{i=1}^S f(v_i) \geq \frac{1}{S} \sum_{i=1}^S f(u_i) = \hat{E}(f(\tilde{u})), \quad \forall \text{ increasing concave function } f(\cdot), \quad (17)$$

where $\hat{E}(\cdot)$ denotes the sample mean. Since the samples taken are independent, the estimator of expected utility in (17) is therefore consistent.

Although the set of increasing concave functions are infinite, the following theorem shows how we could test out this condition efficiently.

Theorem 6.1 (*Levy and Kroll [22]*). *Let \tilde{r}_1 and \tilde{r}_2 be two random variables. Then,*

$$E(f(\tilde{r}_1)) \geq E(f(\tilde{r}_2)), \quad \forall \text{ increasing concave function } f(\cdot)$$

if and only if

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

where the $q_\eta(\tilde{z})$ being the η -quantile defined as follows

$$q_\eta(\tilde{z}) = \inf\{z \mid \Pr(\tilde{z} \leq z) \geq \eta\}.$$

It is a trivial observation from Theorem 6.1 that the condition (17) is equivalent to

$$\Phi_k(v_1, \dots, v_S) \geq \Phi_k(u_1, \dots, u_S), \quad \forall k = 1, \dots, S, \quad (18)$$

where,

$$\Phi_k(y_1, \dots, y_S) = \frac{1}{k} \min_{\substack{W: |W|=k, \\ W \subseteq \{1, \dots, S\}}} \sum_{j \in W} y_j.$$

To see this, let \tilde{v} (respectively \tilde{u}) take discrete values in $\{v_1, \dots, v_S\}$ ($\{u_1, \dots, u_S\}$) with equal probability. In which case, the condition (17) is the same as

$$E[\tilde{v} \mid \tilde{v} \leq q_\eta(\tilde{v})] \geq E[\tilde{u} \mid \tilde{u} \leq q_\eta(\tilde{u})] \quad \forall \eta \in (0, 1).$$

However, since both \tilde{v} and \tilde{u} take discrete values with equal probabilities, we only require to compare the quantiles at $\eta \in \{1/S, 2/S, \dots, 1\}$, in which we will arrive at condition (18).

Figure 2 indicates that

$$\Phi_k(v_1, \dots, v_S) > \Phi_k(u_1, \dots, u_S), \quad \forall k = 1, \dots, S,$$

suggesting that profit profile derived from the risk-neutral LD policy is second order stochastically dominated by the profit profile derived from the risk-neutral FD policy. To further understand the nature of the profit profile derived from the LD policy, we compare, in Figure 3, the parameters of the inventory policy (s_t, S_t) . Interestingly, the replenishment policy under limited demand information tends to over order (higher S_t) in all periods compared with the benchmark risk-neutral FD policy. We, therefore, study numerically whether the mechanism of risk aversion could improve upon the profit profile when the training data is limited.

In Table 4 we compare the risk-averse LD policies under different parameters and the risk-neutral LD policies. Observe that the performance of the risk-neutral LD inventory policy ('risk-neutral') is significantly worse than all the risk-averse LD policies. In particular, with the parameter $\rho = 20$, the sample mean of profit is significantly improved by 4.41%, the standard deviation of profit reduced by 16.5% and the mean loss is reduced by 38.5%.

Figure 4 shows that from sampling point of view, the profit profile based on a risk-neutral LD policy is second order stochastically dominated by the profit profiles associated with the risk-averse LD policy. To understand the improvement over the risk-neutral model, we compare the inventory policy parameters (s_t, S_t) in Figure 5. While there are small changes in the reordering levels, s_t , it is indicative that risk-averse LD policies lead to more conservative ordering quantities which mitigate the risk of over-ordering resulting from limited knowledge of the demand distributions.

This numerical study suggests that when little is known about the demand distributions, the inventory policies derived from a risk-neutral model, the risk-neutral LD policies, could lead to poor profit performance. More importantly, our preliminary computational studies suggests that risk-averse LD policies are also more robust to distributional uncertainties.




Figure 2: Comparison of $\Phi_k(\cdot)$ for risk-neutral policies derived from limited and full demand information.




Figure 3: Plot of (s_t, S_t) against t under limited and full demand information.

ρ	Mean Profit	Standard Deviation of Profit	% Negative Profit	Mean Loss
10	258	215	10.2	71.4
20	261	227	12.0	81.1
40	255	248	15.3	104.6
risk-neutral	250	272	18.0	132.4

Table 4: Summary of results for exponential utility model under limited demand information.




Figure 4: Comparison of $\Phi_k(\cdot)$ for risk-averse policies derived from limited demand information.




Figure 5: Plot of (s_t, S_t) against t under limited demand information.

7 Conclusions

In this paper, we propose a 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. Assuming stationary input parameters, it is natural to expect that a stationary (s, S, p) policy is optimal when price is not a decision variable and a stationary (s, S, A, p) policy is optimal when price is a decision variable. We conjecture that similar to the risk neutral case (see Chen and Simchi-Levi [10]), a stationary (s, S, p) policy is also optimal even when price is a decision variable.
- **Continuous Time Models:** Continuous time models are widely used in finance literature. Thus, it is interesting to extend our periodic review framework to models in which inventory (and pricing) decisions are reviewed in continuous time (similar to Chen and Simchi-Levi [11]) and financial trading takes place in continuous time as well.
- **Portfolio Approach for Supply Contracts:** It is possible 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 [25].
- **The Stochastic Cash Balance Problem:** Recently, Chen and Simchi-Levi [12] 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 3.2, 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.

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. It is possible to extend this framework to different environments, those that go beyond inventory models, for instance, revenue management models.

Throughout this paper, we assume zero lead time. It is well-known that when price is not a decision variable, the structural results of the optimal policy for the risk-neutral inventory models with zero lead time can be extended to risk-neutral inventory models with positive lead time (See Scarf [29]). The idea is to make decisions based on inventory positions, on-hand inventories plus inventory in transit, and reduce the model with positive lead time to one with zero lead time by focusing on the inventory position. To conduct this reduction, we need a critical property that the expectation $E(\cdot)$ of the summation of random variables equals the summation of expectations. Unfortunately, this property does not hold for the certainty equivalent operator $\ln E \exp(\cdot)$ when the random variables are correlated. This implies that a replenishment decision depends not just on inventory positions, but also on the on-hand inventory level and inventories in transit. Thus, our results for risk-averse inventory models with zero lead time cannot be extended to risk-averse inventory models with positive lead time. When price is a decision variable, even under risk-neutral assumption, the structural results for models with zero lead time cannot be extended to models with positive lead time (See Chen and Simchi-Levi [9]).

Finally, we would like to caution the readers about some limitations and practical challenges of the expected utility risk-averse models. Although the expected utility theory is commonly used for modeling risk-averse decision making problems, it does not capture all the aspects of human beings' choice behavior under uncertainty (Rabin [27]). In practice, the set of axioms that the expected utility theory is built upon may be violated. We refer readers to Heyman and Sobel [19] and Fishburn [18] for discussions on the axiomatic game of the expected utility theory. Our model also bears the same practical challenges as other models based on expected utility theory – for example, specifying the decision maker utility function and determining related parameters is not easy. We note that some approaches for assessing decision makers' utility functions were proposed in the decision analysis literature; see, for example, discussions in the textbook by Clemen [14].

Nevertheless, our risk-averse model provides inventory managers an alternative way of making inventory decisions. Our numerical study indicates that the risk-averse models based on the additive exponential utility function are not that sensitive to the choice of ρ . Finally, when there is limited demand information, our numerical study suggests that the solutions obtained from risk-averse models is more robust than the corresponding risk-neutral policy.

A Proof of Theorem 4.1

First consider the last period, period T .

$$\begin{aligned} U_T(x, w) &= \max_{y \geq x, p} E \left[-a_T e^{-(w - k\delta(y-x) + P(y, p; \tilde{\epsilon}_T)) / \rho_T} \right] \\ &= a_T e^{-w / \rho_T} \max_{y \geq x, p} -e^{-k\delta(y-x) / \rho_T} E \left[e^{-P(y, p; \tilde{\epsilon}_T) / \rho_T} \right]. \end{aligned}$$

For simplicity, we don't explicitly write down the constraint $\bar{p}_T \geq p \geq \underline{p}_T$. We follow this convention throughout this appendix.

Define

$$G_T(x) = \max_{y \geq x, p} -k\delta(y-x) - \rho_T \ln \left(E \left[\exp \left(-\frac{1}{\rho_T} P(y, p; \tilde{\epsilon}_T) \right) \right] \right).$$

We have

$$\max_{y \geq x, p} -e^{-k\delta(y-x) / \rho_T} E \left[e^{-P(y, p; \tilde{\epsilon}_T) / \rho_T} \right] = -e^{-G_T(x) / \rho_T}.$$

Thus,

$$U_T(x, w) = -a_T e^{-(G_T(x) + w) / R_T}$$

with R_T defined in Eq. (9).

Now we start induction. Assume

$$U_{t+1}(x, w) = -A_{t+1} e^{-(G_{t+1}(x) + w) / R_{t+1}}.$$

for some constant $A_{t+1} > 0$.

Now we consider period t .

$$U_t(x, w) = \max_{y \geq x, p} E \left[\max_z \left\{ -a_t e^{-\left(w - \frac{1}{1+r_f} z - k\delta(y-x) + P_t(y, p; \tilde{\epsilon}_t) \right) / \rho_t} - A_{t+1} e^{-(G_{t+1}(y-\tilde{d}) + z) / R_{t+1}} \right\} \right],$$

where for simplicity, we use \tilde{d} to denote the demand of period t , which of course is a function of the selling price of this period. For any given (y, p) , the first order optimality condition with respect to z is

$$\frac{1}{\rho_t} a_t e^{-\left(w - \frac{1}{1+r_f} z \right) / \rho_t} e^{(k\delta(y-x) - P_t(y, p; \tilde{\epsilon}_t)) / \rho_t} = \frac{1+r_f}{R_{t+1}} A_{t+1} e^{-z / R_{t+1}} e^{-G_{t+1}(y-\tilde{d}) / R_{t+1}}, \quad (19)$$

or, equivalently (since both a_t and $A_{t+1} > 0$),

$$\ln \frac{a_t}{\rho_t} - \frac{w - z / (1+r_f)}{\rho_t} + \frac{k\delta(y-x) - P_t(y, p; \tilde{\epsilon}_t)}{\rho_t} = \ln \frac{(1+r_f) A_{t+1}}{R_{t+1}} - \frac{z}{R_{t+1}} - \frac{G_{t+1}(y-\tilde{d})}{R_{t+1}}.$$

Thus, for any given (y, p) at state (x, w) and the realization of the current period uncertainty $\tilde{\epsilon}_t$, the optimal banking decision z is

$$z^* = -\frac{\rho_t}{R_t} G_{t+1}(y-\tilde{d}) + \frac{R_{t+1}}{R_t} (-k\delta(y-x) + P_t(y, p; \tilde{\epsilon}_t)) + \frac{R_{t+1}}{R_t} w + \frac{R_{t+1} \rho_t}{R_t} \ln \frac{A_{t+1} (1+r_f) \rho_t}{a_t R_{t+1}},$$

which implies that the optimal consumption decision at time period t is

$$\begin{aligned} f_t^* &= \frac{\rho_t}{R_t} \left[w + (-k\delta(y-x) + P_t(y, p; \tilde{\epsilon}_t)) + \frac{1}{1+r_f} G_{t+1}(y-\tilde{d}) \right] - \frac{R_{t+1}\rho_t}{R_t(1+r_f)} \ln \frac{A_{t+1}(1+r_f)\rho_t}{a_t R_{t+1}} \\ &= \frac{\rho_t}{R_t} \left[w + (-k\delta(y-x) + P_t(y, p; \tilde{\epsilon}_t)) + \frac{1}{1+r_f} G_{t+1}(y-\tilde{d}) \right] + C_t, \end{aligned}$$

if we define constant

$$C_t = -\frac{R_{t+1}\rho_t}{R_t(1+r_f)} \ln \frac{A_{t+1}(1+r_f)\rho_t}{a_t R_{t+1}}.$$

Eq. (19) also implies that

$$\begin{aligned} U_t(x, w) &= -\frac{(1+r_f)R_t}{R_{t+1}} A_{t+1} \max_{y \geq x, p} E \left[-e^{-(z^* + G_{t+1}(y-\tilde{d}))/R_{t+1}} \right] \\ &= -A_t e^{-w/R_t}. \\ &\quad \max_{y \geq x, p} E \left[-\exp \left\{ -\left[G_{t+1}(y-\tilde{d})/(1+r_f) - k\delta(y-x) + P_t(y, p; \tilde{\epsilon}_t) \right] / R_t \right\} \right], \end{aligned}$$

in which

$$A_t = \frac{(1+r_f)R_t}{R_{t+1}} A_{t+1} \left(\frac{A_{t+1}(1+r_f)\rho_t}{a_t R_{t+1}} \right)^{-\frac{R_{t+1}\rho_t}{R_t}} > 0.$$

If we define

$$G_t(x) = \max_{y \geq x, p} -k\delta(y-x) - R_t \ln E \left[\exp \left\{ -\frac{1}{R_t} \left[P_t(y, p; \tilde{d}) + \frac{1}{1+r_f} G_{t+1}(y-\tilde{d}) \right] \right\} \right],$$

we have

$$U_t(x, w) = -A_t \exp \{ -(w + G_t(x)) / R_t \}.$$

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 [29] 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 $k \geq 0$, if for any $x_0 \leq x_1$ and $\lambda \in [0, 1]$,

$$f((1-\lambda)x_0 + \lambda x_1) \leq (1-\lambda)f(x_0) + \lambda f(x_1) + \lambda k. \quad (20)$$

Below we summarize properties of k -convex functions.

Lemma 1 (a) A real-valued convex function is also 0-convex and hence k -convex for all $k \geq 0$. A k_1 -convex function is also a k_2 -convex function for $k_1 \leq k_2$.

(b) If $f_1(y)$ and $f_2(y)$ are k_1 -convex and k_2 -convex respectively, then for $\alpha, \beta \geq 0$, $\alpha f_1(y) + \beta f_2(y)$ 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 $f(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. Let S be a minimum point of g and s be any element of the set

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

Then the following results hold.

- (i) $f(S) + k = f(s) \leq f(y)$, for all $y \leq s$.
- (ii) $f(y)$ is a non-increasing function on $(-\infty, s)$.
- (iii) $f(y) \leq f(z) + k$ for all y, z with $s \leq y \leq z$.

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

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

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

Recently a weaker concept of symmetric k -convexity was introduced by Chen and Simchi-Levi [9, 10] when they analyze the joint inventory and pricing problem with fixed ordering cost.

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

$$f((1-\lambda)x_0 + \lambda x_1) \leq (1-\lambda)f(x_0) + \lambda f(x_1) + \max\{\lambda, 1-\lambda\}k. \quad (21)$$

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 1 and Proposition 2. Finally, notice that the concept of symmetric k -convexity can be easily extended to *multi-dimensional* space.

Lemma 2 (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 $g(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. 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 $\mathfrak{R}^n \times \mathfrak{R}^m \rightarrow \mathfrak{R}$. Assume that for any $x \in \mathfrak{R}^n$ there is a corresponding set $C(x) \subset \mathfrak{R}^m$ such that the set $C \equiv \{(x, y) \mid y \in C(x), x \in \mathfrak{R}^n\}$ is convex in $\mathfrak{R}^n \times \mathfrak{R}^m$. 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 \mathfrak{R}^n .

Proof. For any $x_0, x_1 \in \mathfrak{R}^n$ and $\lambda \in [0, 1]$, let $y_0, y_1 \in \mathfrak{R}^m$ 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. ■

References

- [1] Anvari, M. 1987. Optimality Criteria and Risk in Inventory Models: The Case of the Newsboy Problem. *Journal of Operational Research Society*, 38, 625-632.
- [2] 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.
- [3] Bertsimas, D., A. Thiele. 2003. A Robust Optimization Approach to Supply Chain Management. Working Paper, MIT.
- [4] Bouakiz, M., M. J. Sobel. 1992. Inventory Control with an Exponential Utility Criterion. *Operations Research*, 40, pp. 603-608.
- [5] Caldentey R. and M. Haugh. 2003. *Optimal Control and Hedging of Operations in the Presence of Financial Markets*. working paper.

- [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, MIT.
- [8] Chen, X., M. Sim, D. Simchi-Levi and P. Sun. 2004. Risk Averse Inventory Management. To appear in *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management* Second Edition by D. Simchi-Levi, X. Chen and J. Bramel, Springer, New York.
- [9] 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*.
- [10] Chen, X., D. Simchi-Levi. 2002. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Continuous Review Case. Working paper.
- [11] Chen, X., D. Simchi-Levi. 2003. A New Approach for the Stochastic Cash Balance Problem with Fixed Costs. Working paper, MIT.
- [12] Chen, X., D. Simchi-Levi. 2004. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Infinite Horizon Case. *Mathematics of Operations Research*, **29**(3), pp. 698-723.
- [13] Chung, K. 1990. Risk in Inventory Models: The Case of the Newsboy Problem, Optimality Conditions. *Journal of Operational Research Society*, 41, 173-176.
- [14] Clemen, R. T. 1996. *Making Hard Decisions: An Introduction to Decision Analysis*, Second Edition. PWS-Kent Publishing Co., Boston, MA.
- [15] Eeckhoudt, L., C. Gollier and H. Schlesinger. 1995. The Risk-Averse (and Prudent) Newsboy. *Management Science*, **41**(5), pp. 786-794.
- [16] Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Operations Research*, **47**, No. 3, pp. 454-475.
- [17] Gaur, V. and S. Seshadri. 2004. *Hedging Inventory Risk Through Market Instruments*. working paper.
- [18] Fishburn, P. 1989. Retrospective on the Utility Theory of von Neumann and Morgenstern. *Journal of Risk and Uncertainty*, **2**, pp. 127-158.
- [19] Heyman, D. P., M. J., Sobel. 1982. *Stochastic Models in Operations Research Volume II*. Dover Publications, Inc., Mineola, NY.
- [20] Howard, R. A. 1988. Decision Analysis: Practice and Promise. *Management Science*, 34(6) 679-695
- [21] Lau, H. S. 1980. The Newsboy Problem under Alternative Optimization Objectives. *Journal of the Operations Research Society*, **31**(1), pp. 525-535.

- [22] Levy, H., Y. Kroll. 1978. Ordering uncertain options with borrowing and lending. *Journal of Finance*, **31**(2).
- [23] Keeney, R., H. Raiffa. 1993. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Cambridge University Press.
- [24] Kirkwood, C. W. 2004. Approximating Risk Aversion in Decision Analysis Applications. *Decision Analysis*, 1(1), 51-67
- [25] Martínez-de-Albéniz V., D. Simchi-Levi. 2003. Mean-Variance Trade-offs in Supply Contracts. Working paper, MIT.
- [26] Mas-Collel, A., M. Whinston and J. Green. 1995. *Microeconomic Theory*. Oxford UP.
- [27] Rabin, M. 1998. Psychology and Economics. *Journal of Economics Literature*, **36**(1), pp. 11-46.
- [28] Savage, L. J. 1954. *The Foundations of Statistics*. Wiley, New York.
- [29] 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.
- [30] Schweitzer, M., G. Cachon. 2000. Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence. *Management Science*, **46**(3), 404-420.
- [31] Sethi, P. S. and F. Cheng. 1997. Optimality of (s, S) Policies in Inventory Models with Markovian Demand. *Operations Research*, **45**, No. 6, pp. 931-939.
- [32] Smith, J. E. 1998. Evaluating Income Streams: A Decision Analysis Approach. *Management Science*, **44**(12), 1690-1708.
- [33] Smith, J. E. and R. F. Nau. 1995. Valuing Risky Projects: Option Pricing Theory and Decision Analysis. *Management Science*, **41**(5), 795-816.
- [34] Sobel, M. J. 2003. The Trajan Horse of Expected Discount Utility Theory. Unpublished manuscript, Department of Operations, Case Western Reserve University.
- [35] Song, J. and P. Zipkin. 1993. Inventory Control in a Fluctuating Demand Environment. *Operations Research*. **41**(2), 351-370.