OPERATIONS RESEARCH
Vol. 55, No. 5, September–October 2007, pp. 828–842
ISSN 0030-364X | EISSN 1526-5463 | 07 | 5505 | 0828

informs®

DOI 10.1287/opre.1070.0429

Risk Aversion in Inventory Management

Xin Chen
Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana–Champaign,
Urbana, Illinois 61801, xinchen@uiuc.edu

Melvyn Sim
Singapore-MIT Alliance and NUS Business School, National University of Singapore, Singapore,
dscsimm@nus.edu.sg

David Simchi-Levi
Department of Civil and Environmental Engineering and Engineering System Division, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139, dslevi@mit.edu

Peng Sun
Fuqua School of Business, Duke University, Durham, North Carolina 27708, psun@duke.edu

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 multiperiod inventory models as well as multiperiod models that coordinate inventory and pricing strategies. 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 that the optimal policy is relatively insensitive to small changes in the decision-maker's level of risk aversion.

Subject classifications: inventory/production: policies, uncertainty; decision analysis: risk.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received March 2004; revision received August 2006; accepted August 2006. Published online in *Articles in Advance* September 14, 2007.

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 planners that are insensitive to profit variations.

Evidently, not all inventory planners are risk neutral; many are willing to trade off lower expected profit for downside protection against possible losses. Indeed, experimental evidence suggests that for some products, the so-called *high-profit products*, decision makers exhibit risk-averse behavior; see Schweitzer and Cachon (2000) for more details. Unfortunately, traditional inventory control models fall 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 (1980) analyzes the classical newsvendor model under two different objective functions. In the first objective function, 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 et al. (1995) 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 (2000) analyze the mean-variance trade-offs 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 trade-off of customer waiting time as well as the mean-variance trade-offs of inventory levels. Martínez-de-Albéniz and Simchi-Levi (2003) study the mean-variance trade-offs 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 (1992) is closely related to ours. In this paper, the authors characterize the inventory replenishment strategy so as to minimize the expected utility of the net 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.

Table 1. Summary of previous results and new contributions.

	Price is not a decision		Price is a decision	
	(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 and concave utility	Wealth dependent (modified) base stock	?	Wealth dependent (modified) base stock	?
World-driven model parameter exponential utility	State dependent (modified) base stock	State dependent (s, S)	State dependent (modified) base stock	State dependent (s, S, A, p)
Partially complete financial market exponential utility	State dependent (modified) base stock	State dependent (s, S)	State dependent (modified) base stock	State dependent (s, S, A, p)

*Indicates existing results in the literature.
†Indicates a similar existing result based on a special case of our model.

So far all the papers referenced above assume that demand is exogenous. A rare exception is Agrawal and Seshadri (2000), 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 focuses 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 (1987), which uses the capital asset pricing model (CAPM) to study a newsvendor facing normal demand distribution. Chung (1990) provides an alternative derivation for the result. More recently, Gaur and Seshadri (2005) investigate the impact of financial hedging on the operations decision, and Caldentey and Haugh (2006) 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 multiperiod 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. We assume that the firm we model is a private firm, therefore there is no conflict of interests between share holders and managers. Following Smith (1998), we take the standard economics perspective in which the decision maker maximizes the total expected utility from *consumption* in each time period. In §2, we discuss in more detail the theory of expected utility employed in a multiperiod decision-making framework. We extend our framework in §4 by incorporating a *partially complete* financial market so that the decision maker can hedge operational risk through trading 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 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 (1960) 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 (2004a). 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 nonincreasing 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 1999, Chen and Simchi-Levi 2004a).

Table 1 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 $k = 0$. Note that the optimal policy characterized by Bouakiz and Sobel (1992) 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.

The last row of Table 1 provides information on the optimal policy for a decision maker with an exponential utility function having access to a partially complete financial market. Such a market allows the risk-averse inventory planner to hedge its operational costs and part of the demand risks. 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 §4.

We complement the theoretical results with a numerical study illustrating the effect of risk aversion on the inventory policies.

This paper is organized as follows. In §2, we review classical expected utility approaches in risk-averse valuation. In §3, we propose a model to incorporate risk aversion in a multiperiod inventory (and pricing) setting, and focus on characterizing the optimal inventory policy for a risk-averse decision maker. We then generalize the results in §4 by considering the financial hedging option. Section 5 presents the computational results illustrating the effects of different risk-averse multiperiod inventory models on inventory control policies. Finally, §6 provides some concluding remarks.

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

2. Utility Theory for Risk-Averse Valuations

Modelling 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., Mas-Collel et al. 1995, Chapter 6).

Assume that a decision maker has to make a decision in a single-period problem before uncertainty is resolved. According to 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 over lotteries, 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 1982, Chapters 2–4; Fishburn 1989).

For multiperiod problems, one approach of modeling risk aversion that seems natural is to maximize the expected utility of the net present value of the income cash flow. In calculating the net 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 net present value. Models based on this approach are referred to as the *net present value models*. Sobel (2005) refers to the utility function used in such a framework as the *interperiod utility function*. Note that the net present value models have been employed by Bouakiz and Sobel (1992) and Chen et al. (2004) to analyze the multiperiod 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 1998). 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. More generally a decision maker can trade on financial markets to adjust her consumptions over time.

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

Specifically, 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_t and chooses an operations policy (inventory/pricing) that affects her income cash flow. 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 $w_t + \bar{P}_t$ and decides her consumption level f_t for this period, where $\bar{P}_t$ is the income generated at period t. The remaining wealth, $w_t + \bar{P}_t - f_t$, is then saved (or borrowed, if negative) for the next period. Thus, the next period's initial wealth is

$$w_{t+1} = (1 + r_f)(w_t + \bar{P}_t - f_t).$$

Equivalently, we can model w_{t+1} as a decision variable and calculate the consumption,

$$f_t = w_t - \frac{w_{t+1}}{1 + r_f} + \bar{P}_t.$$

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

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

over the planning horizon $1, \ldots, T$. We call such an approach the *consumption model*. Smith (1998) provides an excellent comparison between the consumption model and the net present value model.

Similar to single-period problems, axiomatic approaches were also employed to derive certain types of utility functions for multiperiod problems (see, e.g., Sobel 2005; Keeney and Raiffa 1993, Chapter 9). In particular, the so-called "additive independence axiom"[1] implies additive utility functions of the following form:

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

That is, the utility of the consumption flow is the summation of the utility from the consumption in each time period, where function u_t is increasing and concave. Sobel (2005) refers to functions u_t as the *intraperiod utility functions*.

As a special case of the general intraperiod utility functions, the exponential utility functions are also commonly used in economics (Mas-Collel et al. 1995) and decision analysis (Smith 1998). 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 (1988) indicates that exponential utility functions are widely applied in decision analysis practice. Kirkwood (2004) shows that in most cases, an appropriately chosen exponential utility function is a very good (local) approximation for general utility functions.

In the next section, we characterize the structures of the optimal inventory policies according to the consumption model. Interestingly, the net present value model is mathematically a special case of the consumption model, as will be illustrated in §3. This implies that the structures of the optimal inventory policies for the consumption models are also valid for the corresponding net present value models.

At this point it is worth mentioning that Savage (1954) unified 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 §4, 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 (1995) and Gaur and Seshadri (2005).

3. Multiperiod 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*

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

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, $t = 1, 2, \ldots,$ 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. Therefore, the inventory level carried over from period t to the next period, x_{t+1}, may be positive or negative. A cost $h_t(x_{t+1})$ is incurred at the end of period t which represents inventory holding cost when $x_{t+1} > 0$ and shortage cost if $x_{t+1} < 0$. For technical reasons, we assume that the function $h_t(x)$ is convex and $\lim_{|x|\to\infty} h_t(x) = \infty$.

At the beginning of period t, the inventory planner 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 $\tilde{\epsilon}_t$, the income at period t is

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

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 that the inventory planner consumes everything, which corresponds to $w_{T+1} = 0$.

According to the consumption model, the inventory planner'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 (or equivalently, the consumption level) for the following optimization problem:

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

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

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

the consumption model reduces to the traditional risk-neutral inventory (and pricing) problem analyzed by Chen and Simchi-Levi (2004a). In this case, we denote $V_t(x)$ to be 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 2004a for more details):

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

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

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

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* $\max_{\bar{p}_t \geqslant p \geqslant \underline{p}_t} g_t(y, p)$ *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 (1960) using the concept of k-convexity; part (b) and part (c) are proved in Chen and Simchi-Levi (2004a) 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 Chen and Simchi-Levi (2004a) 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 model.

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:

$$V_t(x, w) = \max_{y \geqslant x,\, \underline{p}_t \leqslant p \leqslant \bar{p}_t} E_{\tilde{\epsilon}_t}[\overline{W}_t(x, w, y, p;\ \tilde{\epsilon}_t)], \tag{5}$$

in which

$$\overline{W}_t(x, w, y, p;\ \tilde{\epsilon}_t) = \max_{w'} \bigg\{ u_t\bigg(w - \frac{w'}{1+r_f} + \bar{P}_t(x, y, p;\ \tilde{\epsilon}_t)\bigg) + V_{t+1}(y - D_t(p, \tilde{\epsilon}_t), w')\bigg\}, \tag{6}$$

with boundary condition

$$W_T(x, w) = u_T(w + \bar{P}_t(x, y, p;\ \tilde{\epsilon}_t)).$$

Note that unlike the traditional risk-neutral inventory models, where the state variable in the 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 the modified income at period t be

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

The dynamic program (5)–(6) becomes

$$U_t(x, w) = \max_{y \geqslant x,\, \underline{p}_t \leqslant p \leqslant \bar{p}_t} E[W_t(x, w, y, p;\ \tilde{\epsilon}_t)], \tag{7}$$

in which

$$W_t(x, w, y, p;\ \tilde{\epsilon}_t) = \max_{z} \bigg\{ u_t\bigg(w - \frac{z}{1+r_f} - k\delta(y-x) + P_t(y, p;\ \tilde{\epsilon}_t)\bigg) + U_{t+1}(y - D_t(p, \tilde{\epsilon}_t), z)\bigg\}. \tag{8}$$

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, note 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

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

Because $E[W_t(w, y, p;\, \tilde{\epsilon}_t)]$ 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 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 wealth/consumption decisions have any nice structure.

Next, we argue that the net present value model is mathematically a special case of the consumption model. Indeed, if the decision maker's utility functions in each period $t = 1, \ldots, T-1$ are all in the form of $u_t(x) = -\exp(-x/R)$ with $R \to 0^+$, except in period $t = T$, the consumption model (5)–(6) mathematically reduces to the net present value model with intraperiod utility function $U = u_T$. The intuition is also clear. In fact, R (commonly referred to as the "risk-tolerance" parameter) approaching zero implies that the decision maker becomes "extremely risk averse," and thus any negative consumption introduces a negative infinite utility, while any nonnegative consumption introduces zero utility. Therefore, the consumptions in period $t = 1, \ldots, T-1$ have to be nonnegative and the utility is always zero. The decision maker is better off by shifting all the consumptions to the last period, which is equivalent to the net present value model. This also implies that the same structural results in Theorem 3.2 and those to be presented in the next section also hold for the net present value model.

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 reflects the decision maker's attitude toward the utility obtained from different periods.

The beauty of exponential utility functions is that we are able to separately make the inventory decisions without considering the wealth/consumption decisions. This is discovered by Smith (1998) in the decision-tree framework. The next theorem states this result in dynamic programming language. For completeness, a proof is presented in Appendix A.

To state the theorem, we first introduce some notation. For a risk-tolerance parameter R, denote the "certainty equivalent" operator with respect to a random variable $\tilde{\xi}$ to be

$$\mathscr{CE}^R_{\tilde{\xi}}[\tilde{\xi}] = -R \ln E[e^{-\tilde{\xi}/R}].$$

For a decision maker with risk-tolerance R and an exponential utility function, the above certainty equivalent represents the amount of money she feels indifferent to a gamble with random payoff $\tilde{\epsilon}$. Similarly, we denote the "conditional certainty equivalent" operator with respect to a random variable $\tilde{\xi}$ given $\tilde{v}$ to be

$$\mathscr{CE}^R_{\tilde{\xi}|\tilde{v}}[\tilde{\xi}] = -R \ln E[e^{-\tilde{\xi}/R} \mid \tilde{v}].$$

We also consider the "effective risk tolerance" per period, defined as

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

Theorem 3.3. *Assume that $u_t(f) = -a_t e^{-f/\rho_t}$. The inventory decisions in the risk-averse inventory control model*

Equations (5)–(6) *can be calculated through the following dynamic programming recursion*:

$$G_t(x) = \max_{y\geqslant x, \bar{p}_t \geqslant p \geqslant \underline{p}_t} -k\delta(y-x) + \mathscr{C}\mathscr{E}_{\tilde{\epsilon}}^{R_t}\left[P_t(y,p;\tilde{\epsilon}_t) + \frac{1}{1+r_f}G_{t+1}(y - D_t(p,\tilde{\epsilon}_t))\right] \qquad (10)$$

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

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

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

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

The theorem thus implies that when additive exponential utility functions are used: (i) the optimal inventory policy is independent of the wealth level; (ii) the optimal inventory replenishment and pricing decisions can be obtained regardless of the wealth/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 the additive exponential utility function significantly simplifies the problem.

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. 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. Because the net present value model is a special case of the consumption model, our base-stock policy directly implies the result based on the net present value model obtained by Bouakiz and Sobel (1992) using a more complicated argument.

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

Proposition 1. *If a function* $f(x,\tilde{\xi})$ *is concave, k-concave, or symmetric k-concave in x for any realization of* $\tilde{\xi}$, *then for any* $R > 0$, *the function*

$$g(x) = \mathscr{C}\mathscr{E}_{\tilde{\xi}}^{R}[f(x,\tilde{\xi})]$$

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

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

Define $M(x) = E[\exp(f(x-\xi))]$. It suffices to prove that for any x_0, x_1 with $x_0 \leqslant x_1$ and any $\lambda \in [0,1]$,

$$M(x_\lambda) \leqslant M(x_0)^{1-\lambda} M(x_1)^{\lambda}\exp(\lambda K),$$

where $x_\lambda = (1-\lambda)x_0 + \lambda x_1$. Note that

$$\begin{aligned} M(x_\lambda) &\leqslant E[\exp((1-\lambda)f(x_0-\xi) + \lambda f(x_1-\xi) + \lambda K)] \\ &= \exp(\lambda K)E[\exp((1-\lambda)f(x_0-\xi))\exp(\lambda f(x_1-\xi))] \\ &\leqslant \exp(\lambda K)E[\exp(f(x_0-\xi))]^{1-\lambda}E[\exp(f(x_1-\xi))]^{\lambda} \\ &= M(x_0)^{1-\lambda}M(x_1)^{\lambda}\exp(\lambda K), \end{aligned}$$

where the first inequality holds because f is K-convex, and the second inequality follows from the Hölder inequality with $1/p = 1-\lambda$ and $1/q = \lambda$. □

We can now present the optimal policy for the risk-averse multiperiod inventory (and pricing) problem with additive exponential utility functions.

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. See Lemma 1, Proposition 2, Lemma 2, and Proposition 3 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.

3.3. World-Driven Model Parameters

Next, we extend the results for the exponential utility function to the case of "world-driven model parameters." Following Song and Zipkin (1993), 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 depend on the history of the evolution of the business environment. Formally, let finite set Θ_t represent the set of business environments in period t. We use boldface $\mathbf{\Theta}_t = \prod_{\tau=1}^{t} \Theta_\tau$ to represent the set of trajectories of levels from period 1 to t. Each trajectory, $\theta_t \in \mathbf{\Theta}_t$, is referred to as the *state of the world*, which is used to model relevant economic factors that affect the production/inventory cost and revenue. A state of the world uniquely determines the cost parameters and the sufficient statistics of the demand distribution of the inventory model. That is, at each time period t, parameters c_t, h_t, $\underline{p}_t$, and $\bar{p}_t$ are all functions of $\theta_t \in \mathbf{\Theta}_t$ (we express them as $c_t^{\theta_t}$, $h_t^{\theta_t}$, $\underline{p}_t^{\theta_t}$, and $\bar{p}_t^{\theta_t}$, respectively), and the distributions of $\tilde{\alpha}_t$ and $\tilde{\beta}_t$ are also θ_t-dependent. For the state-dependent uncertain demand, we denote $\tilde{\epsilon}_t^{\theta_t} = (\alpha_t^{\theta_t}, \beta_t^{\theta_t})$.

Similarly to what we have done earlier, define

$$\begin{aligned} P_t(y, p, \tilde{\epsilon}_t^{\theta_t};\, \theta_t, \theta_{t+1}) &= \left(\frac{c_{t+1}^{\theta_{t+1}}}{1+r_f} - c_t^{\theta_t}\right) y + \left(p - \frac{c_{t+1}^{\theta_{t+1}}}{1+r_f}\right) D_t(p, \tilde{\epsilon}_t^{\theta_t}) \\ &\quad - h_t\big(y - D_t(p, \tilde{\epsilon}_t^{\theta_t})\big), \end{aligned} \tag{11}$$

which can be thought of as the decision maker's modified income at period t.

The following theorem is a natural extension of Theorems 3.3 and 3.4.

Theorem 3.5. *Separation*: *The optimal inventory and pricing decisions for the world-driven parameter model may be solved through the following dynamic programming recursion*:

$$G_t(x, \theta_t) = c_t^{\theta_t} x + \max_{y, p:\, y \geqslant x,\, \underline{p}^{\theta_t} \leqslant p \leqslant \bar{p}^{\theta_t}} -k\delta(y - x) + L_t(y, p, \theta_t), \tag{12}$$

in which

$$\begin{aligned} L_t(y, p, \theta_t) &= \mathscr{C}\mathscr{E}^{R_t}_{\tilde{\epsilon}^{\theta_t} | \theta_t}\left[\mathscr{C}\mathscr{E}^{R_t}_{\theta_{t+1}|\theta_t}\left[P_t(y, p, \tilde{\epsilon}_t; \theta_t, \theta_{t+1}) + \frac{1}{1+r_f} G_{t+1}\right.\right. \\ &\qquad \left.\left. \cdot\, (y - D_t(p, \tilde{\epsilon}^{\theta_t}), \theta_{t+1})\right]\right] \end{aligned} \tag{13}$$

and with boundary condition $G_{T+1}(x) = 0$. *Thus, the consumption decisions are decoupled from the inventory* (*pricing*) *decisions.*

Structural policy: *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* $G_t(x, \theta_t)$ *are* k_t*-concave in* x *and a* θ_t*-dependent* (s, S) *inventory policy is optimal.*

(b) *For the general case,* $G_t(x, \theta_t)$ *are symmetric* k_t*-concave in* x *for any given* θ_t *and a* θ_t*-dependent* (s, S, A, p) *policy is optimal.*

Note that the separation result in Theorem 3.5 could be extended to the situation where the fixed cost k is world driven. If we further have the following condition for all θ_t,

$$(1 + r_f)k^{\theta_t} \geqslant \max_{\theta_{t+1}:\, \theta_t \in \theta_{t+1}} k^{\theta_{t+1}},$$

we have that the value functions $G_t(x, \theta)$ are $k_t^{\theta_t}$-concave (symmetric $k_t^{\theta_t}$-concave) and a θ_t-dependent (s, S) (θ_t-dependent (s, S, A, p)) inventory policy is optimal.

In the next section, we further extend the world-driven parameter model by considering the situation that the inventory planner has access to a financial market to hedge the risks associated with fluctuations in the states of the world.

4. Multiperiod Inventory Models with Financial Hedging Opportunities

The modern financial market provides opportunities to replicate many of the changes in the state of the world. Therefore, a risk-averse inventory planner may use the financial market to hedge the risks from changes in the business environment. For example, if the production cost is a function of the oil price, the inventory planner may hedge the oil price risks through trading financial securities on oil prices. Similarly, if the demand distribution is affected by the general economic situation, financial instruments on the market indices provide the possibility of hedging the risks of general trend in demand. In this section, we extend our previous framework by assuming that the decision maker has opportunities of hedging operational risk through trading financial securities in a so-called "partially complete" financial market.

Similarly to the previous section, we consider a risk-averse inventory planner who has to make replenishment (and pricing) decisions over a finite time horizon with T periods. The inventory, pricing, and trading decisions are made at time periods $t = 1, \ldots, T$. The model parameters are "world driven" as defined in the last section. In this section, we explicitly assume that the fixed cost k is world driven. Besides the risk-free borrowing and saving opportunity (cash), we assume that there are another N financial securities in the financial market. To simplify notation and analysis, assume that these securities do not pay dividends during the time horizon $1, \ldots, T$. 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). We follow the usual assumption in the real-options literature that the financial security could be traded at the exact desired amount and there is no transaction fee.

Following Smith and Nau (1995), we refer to the risks associated with the evolution of the state of the world θ_t as *market risk*—it can be fully hedged in the financial market. On the other hand, given the state of the world θ_t, the demand uncertainty in our model is the so-called *private risk* that cannot be hedged in the financial market. The existence of the private risk contributes to the incompleteness of the financial market, which Smith and Nau (1995) called a *partially complete* financial market. Note 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 in the following subsection. The next subsection is devoted to the formal description of a *complete* financial market. The discussion of the complete market and no arbitrage conditions are standard in the finance literature. For a discrete-time treatment of such a financial market, we refer readers to Pliska (1997).

4.1. Complete Financial Market and Risk-Neutral Probabilities

Formally, a financial market is "complete" if we have the following conditions.

Assumption 2. (1) *Q_{it} only depends on the state of the world $\theta_t \in \mathbf{\Theta}_t$. That is, for any trajectory $\theta_T = \{\vartheta_1, \ldots, \vartheta_T\}$ and its subtrajectories $\theta_t = \{\vartheta_1, \ldots, \vartheta_t\}$, 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 $\{\mathcal{P}_t(\theta_t)\}_{\forall \theta_t \in \mathbf{\Theta}_t}$ is a linear combination of the vectors $\mathbf{1}$ and $\{Q_{1t}(\theta_t)\}_{\forall \theta_t \in \mathbf{\Theta}_t}, \ldots, \{Q_{Nt}(\theta_t)\}_{\forall \theta_t \in \mathbf{\Theta}_t}$. Here, $\mathcal{P}_t(\cdot)$ is any mapping from $\mathbf{\Theta}_t$ to a real number representing a state of the world adapted cash flow.*

(3) *Disclosed demand information in each time period is not correlated with 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.*

Intuitively speaking, arbitrage free means that one cannot guarantee positive gain only through trading financial securities on the market.

Formally, to define arbitrage opportunities, we need to introduce the notion of a *self-financing trading strategy*, a well-known concept in finance. A self-financing trading strategy is an $N+1$-dimensional vector of θ_t adapted stochastic processes $\{(w_t, \mathbf{w}_t)\}_{t=1,\ldots,T}$ such that

$$(1+r_f)w_t(\theta_{t-1}) + Q_{\cdot t}(\theta_t)^\top \mathbf{w}_t(\theta_{t-1}) = (1+r_f)w_{t+1}(\theta_t) + Q_{\cdot t}(\theta_t)^\top \mathbf{w}_{t+1}(\theta_t) \qquad (14)$$

for each time period t and for any state-of-the-world trajectories θ_{t-1} and θ_t such that θ_{t-1} is a subtrajectory of θ_t. To be specific, $(w_t, \mathbf{w}_t)$ represents the positions of cash and risky financial securities at the beginning of period t. That is, the number of shares in security i is $\mathbf{w}_t(i)$. Note that $(w_t, \mathbf{w}_t)$ is determined through the trading in period $t-1$ based on the information θ_{t-1} that was available. Equation (14) implies that the values of the portfolio before and after the financial trading in period t are the same. Therefore, no money is added to or subtracted from the portfolio throughout the planning horizon according to a self-financing trading strategy. With the help of the notion of self-financing trading strategies, the arbitrage-free condition can be represented as the following: there does not exist a self-financing trading strategy $\{(w_t, \mathbf{w}_t)\}_{t=1,\ldots,T}$ such that

$$w_1 + \mathbf{w}_1^\top Q_{\cdot 1} = 0,$$
$$w_T(\theta_{T-1}) + \mathbf{w}_T(\theta_{T-1})^\top Q_{\cdot T}(\theta_T) \geqslant 0 \quad \forall\, \theta_T, \quad \text{and}$$
$$w_T(\theta_{T-1}) + \mathbf{w}_T(\theta_{T-1})^\top Q_{\cdot T}(\theta_T) > 0 \quad \text{for some } \theta_T.$$

Assumption 2, parts (1) and (2), also imply 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 $\mathbf{\Theta}$ such that for all $t = 1, \ldots, T$,

$$Q_{it-1}(\theta_{t-1}) = \sum_{\theta_t} \frac{1}{1+r_f} \pi(\theta_t \mid \theta_{t-1}) Q_{it}(\theta_t), \qquad (15)$$

in which $\pi(\theta_t \mid \theta_{t-1})$ is the risk-neutral probability of observing the trajectory θ_t given the subtrajectory up to time period t is θ_{t-1}.

In the sequel, we use $E_\pi[\cdot \mid \theta_t]$ to denote the conditional expectation taken with respect to the risk-neutral probability distribution π, while $E_{\theta_{t+1}}[\cdot \mid \theta_t]$ 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 the notation $E_{\tilde{\epsilon}_t^{\theta_t}}[\cdot]$. Therefore, Equation (15) can be equivalently expressed as

$$Q_{it-1}(\theta_{t-1}) = E_\pi[Q_{it}(\theta_t)/(1+r_f)].$$

As was pointed out by one of the referees, our model can be extended to the case when the risk-free borrowing and saving interest rate r_f is world driven as well. In a complete financial market, a nonstate driven interest rate (in terms of dollars) exists anyway, which will be used to serve our model if the utility functions are in terms of payoff in dollars. As a matter of fact, in a complete financial market, we can design a portfolio such that one dollar worth of such a portfolio is always worth some fixed amount $\sigma_t > 0$ in any given period t regardless of the the state of the world in period t. Therefore, $\sigma_t/\sigma_{t-1} - 1$ could be considered as the risk-free interest rate for time period t. For simplicity of exposition, we assume that r_f is the same across different time periods.

4.2. Partially Complete Financial Market

Following the notations introduced before, we use the N-dimensional vector $\mathbf{w}_t$ to express the inventory planner's financial market position at time period t and the scalar w_t to represent the amount of cash in the bank at period t. In the beginning of each time period t, the decision maker observes the current state of the world θ_t, the inventory level x_t, and the financial market position $(w_t, \mathbf{w}_t)$, and then makes the inventory and pricing decisions y_t and p_t. After observing the realized demand (and thus the income cash flow $\bar{P}_t(x_t, y_t, p_t, \tilde{\epsilon}_t^{\theta_t}; \theta_t)$), she makes the decision on the next period market position $(w_{t+1}, \mathbf{w}_{t+1})$ 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

$$w_{t+1} = (1+r_f)\big(w_t + \bar{P}_t(x_t, y_t, p_t, \tilde{\epsilon}_t^{\theta_t}; \theta_t) + (\mathbf{w}_t - \mathbf{w}_{t+1})^\top Q_{\cdot t} - f_t\big).$$

Equivalently, we have

$$f_t\big(w_t, w_{t+1}, \mathbf{w}_t, \mathbf{w}_{t+1}, x_t, y_t, p_t, \tilde{\epsilon}_t^{\theta_t}; \theta_t\big) = \big(\mathbf{w}_t - \mathbf{w}_{t+1}\big)^\top Q_{\cdot t} + \bar{P}_t\big(x_t, y_t, p_t, \tilde{\epsilon}_t^{\theta_t}; \theta_t\big) + w_t - \frac{w_{t+1}}{1+r_f}.$$

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

$$V_t(x, w, \mathbf{w}, \theta_t) = \max_{y,\, p:\, y\geqslant x,\ \underline{p}_t^{\theta_t} \leqslant p \leqslant \bar{p}_t^{\theta_t}} E_{\tilde{\epsilon}_t^{\theta_t}}\big[\overline{W}_t(x, w, \mathbf{w}, y, p;\ \tilde{\epsilon}_t^{\theta_t}, \theta_t)\big], \quad (16)$$

in which

$$\overline{W}_t(x, w, \mathbf{w}, y, p; \tilde{\epsilon}_t^{\theta_t}, \theta_t) = \max_{z,\mathbf{z}}\Big\{u_t\big(f_t(w, z, \mathbf{w}, \mathbf{z}, x, y, p, \tilde{\epsilon}_t^{\theta_t};\ \theta_t)\big) + E_{\theta_{t+1}}\big[V_{t+1}(y - D_t(p, \tilde{\epsilon}_t^{\theta_t}), z, \mathbf{z}, \theta_{t+1}) \mid \theta_t\big]\Big\}, \quad (17)$$

with boundary condition

$$\overline{W}_T\big(x, w, \mathbf{w}, y, p;\ \tilde{\epsilon}_T^{\theta_T}, \theta_T\big) = u_T\big(w_t + \mathbf{w}_t^\top Q_{\cdot T}(\theta_T) + \bar{P}_T(x, y, p, \tilde{\epsilon}_T^{\theta_T};\ \theta_T)\big).$$

Note 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 (1995), we know that a risk-averse inventory planner 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 §5 of Smith and Nau (1995) and the previous section of this paper. Define the modified income flow in period t, $P_t(y, p, \tilde{\epsilon}_t^{\theta_t};\ \theta_t, \theta_{t+1})$, as in Equation (11).

THEOREM 4.1. *Separation*: *The inventory and pricing decisions in the risk-averse inventory model with financial hedging Equations* (16)–(17) *can be calculated through the following dynamic programming recursion*:

$$G_t(x, \theta_t) = c_t^{\theta_t} x + \max_{y,\, p:\, y\geqslant x,\ \underline{p}^{\theta_t} \leqslant p \leqslant \bar{p}^{\theta_t}} -k_t^{\theta_t}\delta(y-x) + L_t(y, p, \theta_t), \quad (18)$$

in which

$$L_t(y, p, \theta_t) = \mathscr{C}\mathscr{E}^{R_t}_{\epsilon_t^{\theta_t}|\theta_t}\bigg[E_\pi\bigg[P_t(y, p, \tilde{\epsilon}_t^{\theta_t}; \theta_t, \theta_{t+1}) + \frac{1}{1+r_f}G_{t+1}(y - D_t(p, \tilde{\epsilon}^{\theta_t}), \theta_{t+1})\,\Big|\,\theta_t\bigg]\bigg] \quad (19)$$

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

Structural policy: *If, in addition,* $k_t^{\theta_t} \geqslant E_\pi[k_{t+1}^{\theta_{t+1}} \mid \theta_t]$, *then the following structural results for the optimal inventory* (*pricing*) *policies hold.*

(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 a* 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 a* 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 (17), the expectation operator $E_{\theta_{t+1}}[\cdot \mid \theta_t]$ 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, 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 (1995) for the detailed description of such an optimal trading strategy.[2]

It is also interesting to compare the dynamic program for the financial hedging case (18)–(19) with the one without the financial hedging opportunity in (12)–(13). The only difference in the expressions is that the certainty equivalent operator with respect to the state-of-the-world transitions in (13) is replaced by an expectation operator with respect to the risk-neutral probability.

It is appropriate to point out that the restriction on fixed costs in the theorem is similar to the assumptions made in Sethi and Cheng (1997) 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 a numerical study. We consider an additive exponential utility model in which $\rho_t = \rho$ for all $t = 1, \ldots, T$. Assuming 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 with many different demand distributions and inventory scenarios and observed similar trends in profit profile and changes in the inventory policy 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 than or equal to y. Because the demand distribution is bounded and discrete, we can easily evaluate expectations within the dynamic programming recursion and compute the optimal policy exactly.

Table 2. Parameters of the inventory model.

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

To evaluate the inventory policies derived, we analyze the inventory policies 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 Tth 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 $S = 10{,}000$. For each risk parameter $\rho \in \{10, 20, 40\}$, we construct the optimal risk-averse inventory policy.

We now study numerically how the replenishment policies change as we vary the risk-aversion level. That is, because 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 nine 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 a monotonicity property is much desired, we have numerical examples that violate this property as we change the risk-aversion level.

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

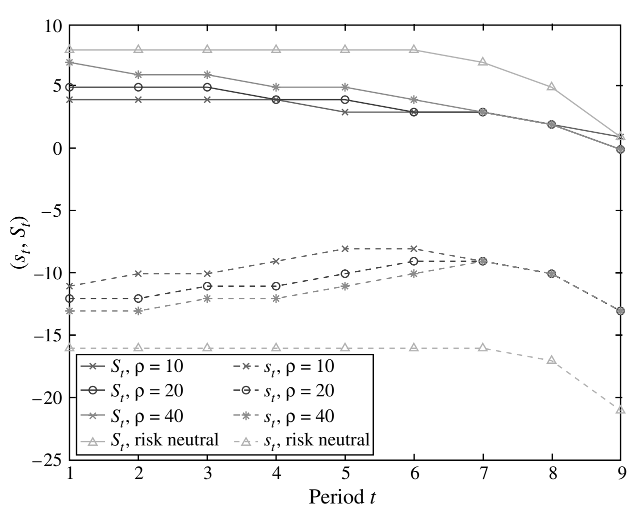

To test the sensitivity of the parameters of the optimal policy to changes in the level of risk aversion, we track the changes in the parameters of the optimal policy as we gradually increase the parameter ρ. It is interesting to observe that while the risk-averse and risk-neutral policies are different, the policy changes resulting from small changes in the risk-tolerance level are quite small. For instance, the optimal policy remains the same as we vary $\rho = 15, 16, \ldots, 24$. Therefore, we conclude (numerically) that the optimal policy is relatively insensitive to small changes in the decision-maker's level of risk aversion.

6. 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) 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 similarly to the risk-neutral case (see Chen and Simchi-Levi 2004b), 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 the 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 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 multiperiod framework. Observe that a different risk-averse model, based on the mean-variance trade-off in supply contracts, cannot be easily extended to a multiperiod framework, as pointed out by Martínez-de-Albéniz and Simchi-Levi (2006).
- **The Stochastic Cash-Balance Problem:** Recently, Chen and Simchi-Levi (2003) 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, similarly to what we did in §3.2, that this type of policy 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 multiechelon inventory models. In addition, it may be possible to extend this framework to different environments that go beyond inventory models (for example, revenue management models).

Another possible extension is to include positive lead time. Indeed, 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 1960). 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 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 assumptions, the structural results for models with zero lead time cannot be extended to models with positive lead time (see Chen and Simchi-Levi 2004b).

Finally, we would like to caution the readers about some limitations and practical challenges of our model. First, the assumption that the savings and borrowing rates are identical may not hold in practice, especially for the majority of manufacturing firms, where the borrowing rate is typically higher than the savings rate. Similar to many economic and financial models, our results depend on this assumption.

Second, although 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 1998). In practice, the set of axioms that expected utility theory is built upon may be violated. We refer readers to Heyman and Sobel (1982) and Fishburn (1989) for discussions on the axiomatic game of 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 are not easy. We note that some approaches for assessing the decision-makers' utility functions were proposed in the decision analysis literature; see, for example, discussions in the textbook by Clemen (1996).

Nevertheless, our risk-averse model provides inventory planners 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.

Appendix A. Proof of Theorem 3.3

First, consider the last period, period T.

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

For simplicity, we do not explicitly write down the constraint $\bar{p}_T \geqslant p \geqslant \underline{p}_T$. We follow this convention throughout this appendix.

Define

$$G_T(x) = \max_{y\geqslant x,p} -k\delta(y-x) + \mathscr{C}\mathscr{E}^{\rho_T}_{\tilde{\epsilon}_T}[P(y,p;\,\tilde{\epsilon}_T)].$$

We have

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

Thus,

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

with R_T defined in Equation (9).

Now we start induction. Assume that

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

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

Now we consider period t:

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

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

$$\frac{1}{\rho_t} a_t e^{-(w-(1/(1+r_f))z)/\rho_t} e^{(k\delta(y-x)-P_t(y,p;\,\tilde{\epsilon}_t))/\rho_t}$$
$$= \frac{1+r_f}{R_{t+1}} A_{t+1} e^{-z/R_{t+1}} e^{-G_{t+1}(y-\tilde{d})/R_{t+1}}, \tag{A1}$$

or, equivalently (because 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

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

if we define constant

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

Equation (A1) also implies that

$$U_t(x,w) = -\frac{(1+r_f)R_t}{R_{t+1}} A_{t+1} \max_{y\geqslant x,p} E\big[-e^{-(z^*+G_{t+1}(y-\tilde{d}))/R_{t+1}}\big]$$
$$= -A_t e^{-w/R_t}$$
$$\cdot \max_{y\geqslant x,p} E[-\exp\{-[G_{t+1}(y-\tilde{d})/(1+r_f) - k\delta(y-x) + P_t(y,p;\tilde{\epsilon}_t)]/R_t]\},$$

in which

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

If we define

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

we have

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

Appendix 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 2003 for more details).

The concept of k-convexity was introduced by Scarf (1960) to prove the optimality of an (s,S) inventory policy for the traditional inventory control problem.

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

$$f((1-\lambda)x_0 + \lambda x_1) \leqslant (1-\lambda)f(x_0) + \lambda f(x_1) + \lambda k. \quad \text{(B1)}$$

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 $k \geqslant 0$. A k_1-convex function is also a k_2-convex function for $k_1 \leqslant k_2$.*

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

$$\{x \mid x \leqslant S, f(x) = gf(S) + k\}.$$

Then, the following results hold:

(i) *$f(S) + k = f(s) \leqslant f(y)$ for all $y \leqslant s$.*

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

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

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

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

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

Recently, a weaker concept of symmetric k-convexity was introduced by Chen and Simchi-Levi (2002a, b) when they analyzed the joint inventory and pricing problem with fixed-ordering cost.

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

$$f((1-\lambda)x_0 + \lambda x_1) \leqslant (1-\lambda)f(x_0) + \lambda f(x_1) + \max\{\lambda, 1-\lambda\}k. \quad \text{(B2)}$$

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, note that the concept of symmetric k-convexity can be easily extended to multidimensional space.

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

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

$$X := \{x \mid x \leqslant S, g(x) = g(S) + k \text{ and } g(x') \geqslant g(x) \text{ for any } x' \leqslant x\}.$$

Then, we have the following results:

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

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

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

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

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

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

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,

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

Therefore, g is symmetric K-convex. □

Endnotes

1. "Attributes $X_1, X_2, \ldots, X_n$ are *additive independent* if preferences over lotteries on $X_1, X_2, \ldots, X_n$ depend only on their marginal probability distributions and not on their joint probability distribution." (See Keeney and Raiffa 1993, Chapter 6.5, p. 295.) Note that the above definition is in the multiattribute preference setting. Preferences over money at different points of time could be treated as multiattribute preferences.

2. We caution reader on the difference in the notation of this paper and Smith and Nau (1995). In this paper, we measure the prices of financial securities in the period t dollar, while Smith and Nau measure in period 1 dollar.

Acknowledgments

The authors gratefully acknowledge many valuable suggestions made by James E. Smith, the associate editor, and two anonymous referees, that led to significant improvement of this paper. They also thank Ralph Keeney for pointing out some relevant references. This research was supported in part by the Singapore-MIT Alliance, ONR contracts N00014-95-1-0232 and N00014-01-1-0146, NSF contracts DMI-9732795, DMI-0085683, and DMI-0245352, and NUS academic research grant R-314-000-066-122.

References

Agrawal, V., S. Seshadri. 2000. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing Service Oper. Management* **2**(4) 410–423.

Anvari, M. 1987. Optimality criteria and risk in inventory models: The case of the newsboy problem. *J. Oper. Res. Soc.* **38** 625–632.

Bouakiz, M., M. J. Sobel. 1992. Inventory control with an exponential utility criterion. *Oper. Res.* **40** 603–608.

Caldentey, R., M. Haugh. 2006. Optimal control and hedging of operations in the presence of financial markets. *Math. Oper. Res.* **31**(2) 285–304.

Chen, F., A. Federgruen. 2000. Mean-variance analysis of basic inventory models. Working paper, Columbia University, New York.

Chen, X. 2003. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Chen, X., D. Simchi-Levi. 2003. A new approach for the stochastic cash balance problem with fixed costs. Working paper, Massachusetts Institute of Technology, Cambridge, MA.

Chen, X., D. Simchi-Levi. 2004a. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Oper. Res.* **52**(6) 887–896.

Chen, X., D. Simchi-Levi. 2004b. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Math. Oper. Res.* **29**(3) 698–723.

Chen, X., M. Sim, D. Simchi-Levi, P. Sun. 2004. Risk-averse inventory management. D. Simchi-Levi, X. Chen, J. Bramel, eds. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management*, 2nd ed. Springer, New York, 158–163.

Chung, K. 1990. Risk in inventory models: The case of the newsboy problem, optimality conditions. *J. Oper. Res. Soc.* **41** 173–176.

Clemen, R. T. 1996. *Making Hard Decisions: An Introduction to Decision Analysis*, 2nd ed. PWS-Kent Publishing Co., Boston, MA.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Sci.* **41**(5) 786–794.

Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3) 454–475.

Fishburn, P. 1989. Retrospective on the utility theory of von Neumann and Morgenstern. *J. Risk Uncertainty* **2** 127–158.

Gaur, V., S. Seshadri. 2005. Hedging inventory risk through market instruments. *Manufacturing Service Oper. Management* **7**(2) 103–120.

Heyman, D. P., M. J. Sobel. 1982. *Stochastic Models in Operations Research*, Vol. II. Dover Publications, Inc., Mineola, NY.

Howard, R. A. 1988. Decision analysis: Practice and promise. *Management Sci.* **34**(6) 679–695.

Keeney, R., H. Raiffa. 1993. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Cambridge University Press, Cambridge, UK.

Kirkwood, C. W. 2004. Approximating risk aversion in decision analysis applications. *Decision Anal.* **1**(1) 51–67.

Lau, H. S. 1980. The newsboy problem under alternative optimization objectives. *J. Oper. Res. Soc.* **31**(6) 525–535.

Martínez-de-Albéniz, V., D. Simchi-Levi. 2003. Mean-variance trade-offs in supply contracts. *Naval Res. Logist.* **53**(7) 603–616.

Mas-Collel, A., M. Whinston, J. Green. 1995. *Microeconomic Theory*. Oxford University Press, Oxford, UK.

Pliska, S. R. 1997. *Introduction to Mathematical Finance: Discrete Time Models*. Blackwell Publishers Ltd., Oxford, UK.

Rabin, M. 1998. Psychology and economics. *J. Econom. Literature* **36**(1) 11–46.

Savage, L. J. 1954. *The Foundations of Statistics*. Wiley, New York.

Scarf, H. 1960. The optimality of (s, S) policies for the dynamic inventory problem. *Proc. 1st Stanford Sympos. Math. Methods Soc. Sci.* Stanford University Press, Stanford, CA, 196–202.

Schweitzer, M., G. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Sci.* **46**(3) 404–420.

Sethi, P. S., F. Cheng. 1997. Optimality of (s, S) policies in inventory models with Markovian demand. *Oper. Res.* **45**(6) 931–939.

Smith, J. E. 1998. Evaluating income streams: A decision analysis approach. *Management Sci.* **44**(12) 1690–1708.

Smith, J. E., R. F. Nau. 1995. Valuing risky projects: Option pricing theory and decision analysis. *Management Sci.* **41**(5) 795–816.

Sobel, M. J. 2005. Discounting and risk neutrality. Unpublished manuscript, Department of Operations, Case Western Reserve University, Cleveland, OH.

Song, J., P. Zipkin. 1993. Inventory control in a fluctuating demand environment. *Oper. Res.* **41**(2) 351–370.