New Sufficient Conditions for (s, S) Policies to be Optimal in Systems with Multiple Uncertainties

Lucy Gongtao Chen*, Lawrence W. Robinson**,
Robin O. Roundy*** [1], Rachel Q. Zhang **** [2]

** Department of Decision Sciences, NUS Business School, National University of Singapore, Singapore, 117592*

*** Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853*

**** Department of Mathematics, Brigham Young University, Provo, UT 84602*

***** Department of Industrial Engineering and Logistics Management Hong Kong University of Science and Technology Clear Water Bay, Kowloon, Hong Kong*

bizcg@nus.edu.sg, lwr2@cornell.edu**, robin@math.byu.edu ***, rzhang@ust.hk*****

Abstract

In today's business environment, unpredictable economic and non-economic forces can affect firms' operational costs and discount factors, as well as demand. In this paper, we incorporate these uncertainties into a single-product, periodic-review, finite-horizon stochastic inventory system by modeling operational costs, discount factors, and demands as stochastic processes that evolve over time. We study three stockout protocols and establish conditions under which (s, S) inventory policies are optimal when discount factors, operational costs, and demands are stochastic and correlated both to one another and over time. Examples are provided to demonstrate non-trivial optimal policies in the absence of these sufficient conditions.

1 Introduction

Virtually all research on inventory management to date has assumed that operational costs as well as discount factors are fixed and known in advance. Under these assumptions, it is well established that variants of (Q, r) and (s, S) policies are generally optimal (Scarf 1959, Veinott 1966). However, it has long been recognized that market risks can significantly affect a firm's operational costs as well as its discount factor. For example, the purchasing price can fluctuate with the exchange rate if the product is sourced from a foreign country, and the discount factor can change with the riskiness of the market in which the business is conducted. Furthermore, customer demand can be highly correlated with general market conditions (Gaur and Seshadri 2005) or (for certain products such as skis or air conditioners) the weather. In order to fully capture the dynamic business environment

[1]Research supported partially by NSF grants DMI-0075627 and DMI-0500263, and the Querétaro Campus of the Instituto Tecnológico y de Estudios Superiores de Monterrey.

[2]Research supported partially by the Hong Kong Research Grants Council Grants No. 622510 and 622811.

and to understand its implications for firms' operations, it is necessary to treat operational cost parameters (*e.g.*, inventory holding costs, stockout costs, and purchasing prices), discount factors, and demands as stochastic processes that are influenced by the state of the world and that evolve over time. The state of the world can include both economic variables (such as interest rate and exchange rate) and non-economic variables (such as weather).

Although the literature on inventory management with multiple sources of uncertainty and volatility is scarce, one line of research examines the non-stationarity of the demand. Gaur and Seshadri (2005) consider a newsvendor problem where the demand is correlated with the price of a financial security. They study how to use financial options to construct hedging transactions to minimize profit variance. Graves (1999) considers a system where the demand process is an integrated moving average process. Johnson and Thompson (1975) study a discrete-time model where demand is either an autoregressive or a moving average process, and show that the myopic policy is sometimes optimal. Iglehart and Karlin (1962) provide optimality algorithms for the case where the demand follows a Markov process. In a similar but more general problem setting, Song and Zipkin (1993) study the case when the demand is governed by the state of the world and follows a Markov process. They show that a state-dependent (r, S) policy is optimal, where r is the reorder point and S is the order-up-to level. Similar inventory models have also been studied by Karlin and Fabens (1959), Sethi and Cheng (1997), and Cheng and Sethi (1999), where the former consider the problem complex and optimize over a restricted class of ordering policies; the latter two generalize the cost functions and optimize over the class of all history-dependent policies to show that a state-dependent (s,S) policy is optimal.

While the aforementioned literature focuses on time-varying stochastic demand, Fabian *et al.* (1959) examine a raw material inventory problem in a periodic setting where the unit purchasing price varies over time. Building on this work, Kalymon (1971) studies an inventory system where the unit purchasing price follows a Markov process and demand is determined by the price. Assuming the holding and shortage cost function is convex, he shows that a state-dependent (s,S) policy will be optimal.

To understand the optimal inventory replenishment strategy in an unstable environment, we incorporate multiple sources of uncertainty into a single-product, periodic-review, finite-horizon stochastic inventory system by allowing the stochastic discount factors and operational costs to change over time, based on the evolving state of the world. We also allow (but do not require) demand to be correlated with the state of the world. In the presence of multiple sources of uncertainty within the business environment, it is unclear *a priori* what type of inventory policy a firm should follow. Our paper provides comprehensive analyses for three different stockout protocols (pay-to-order and pay-to-delivery under backlogging, and lost sales) and develop conditions under which state-dependent $(s,\ S)$ inventory policies are optimal when discount factors, operational costs, and demands are stochastic and correlated. These conditions result from the impact that the state of the world has on the inventory holding and backorder costs, and may fail in turbulent business environment. Our results enable firms to better manage their inventory in an unstable business environment, and so improve their profitability. To the best of our knowledge, ours is the first inventory model that incorporates stochastic discount factors, operational costs, and demands that all evolve in response to changing business conditions, although the latter two have been studied in literature with a convexity assumption imposed on the holding and shortage cost function.

The paper is organized as follows. In Section 2 we introduce an inventory model with stochastic discount factors and cost parameters. In Section 3 we establish conditions for the optimality of the (s, S) policy, and discuss the implications of these conditions. In Section 4 we provide examples of non-trivial policies when the conditions guaranteeing the optimality of (s, S) policies fail. We summarize in Section 5.

2 Problem Formulation

2.1 Model Description

We consider a periodic review inventory system with a single product. At the beginning of each period, purchasing decisions and payments to the suppliers are made. At the end of each period, demand occurs and available inventory is delivered to customers. Holding and stockout costs are incurred, and revenues are realized.

We consider both the backorder case and the lost sales case (referred to as LS). Under backorder, we consider two cases: the "pay-to-order" (PTO) case where the customers pay when they place their orders, and the "pay-to-delivery" (PTD) case where they pay when their orders are delivered. In both cases, we assume that corporate policy calls for backordered demand to be met as soon as inventory becomes available.

The following notation will be used in our model. We use boldfaced letters to represent vectors whose dimensions will be clear from the context.

T = the length of the planning horizon,

$$t = \text{the period index, increasing over time,}$$

$$L = \text{replenishment lead time (assumed to be zero in the LS case),}$$

$$x_t = \text{inventory position at the beginning of period } t\text{, before ordering,}$$

$$y_t = \text{inventory position at the beginning of period } t\text{, after ordering,}$$

$$D_t = \text{demand during period } t,$$

$$D_{t,w} = \sum_{j=t}^{w} D_j, \text{total demand from period } t \text{ through period } w,$$

$$\mathbf{D}_{t:w} = (D_t, D_{t+1}, ..., D_w), \text{ vector of demands from periods } t \text{ through } w,$$

$$\mathbf{Z}_t = \text{vector representing the state of the world as of the beginning of period } t,$$

$$\mathbf{z}_t = \text{a realization of } \mathbf{Z}_t,$$

$$\mathbf{Z}_{t:w} = (\mathbf{Z}_t, \cdots, \mathbf{Z}_w), \text{set of vectors for the state of the world in periods } t \text{ through } w.$$

Note that we restrict our attention to the case of zero lead time under Lost Sales. This is because, with a positive lead time, it is unclear what the optimal policy is even under the traditional inventory model, due to curse of dimensionality. As such, we focus on the scenario where (s, S) policies are known to be optimal under the traditional setting.

The State-of-the-World Vector The vector that describes the state of the world, $\mathbf{Z}_t$, may include exchange rates, stock market indices, interest rates, weather conditions, *etc.* — all of which evolve over time. We assume that $\{\mathbf{Z}_t,\ 1 \leq t \leq T+1\}$ is a Markov chain, possibly with an infinite number of states, and possibly time-varying.

In this paper, t will usually refer to the current time period, so that $\mathbf{z}_\tau$ has already been observed for $\tau \leq t$. The distributions of $\mathbf{Z}_{t+1:T+1} = (\mathbf{Z}_{t+1}, \mathbf{Z}_{t+2}, ..., \mathbf{Z}_{T+1})$ and $\mathbf{D}_{t:T} = (D_t, D_{t+1}, ..., D_T)$ may be conditioned on the current (observed) vector $\mathbf{z}_t$. We allow, but do not require, demand to be correlated with the state of the world.

Stochastic Discount Factors We define $\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1})$ to be the discount factor that represents the value, at the beginning of period t, of a unit cash flow at the beginning of period $t+1$. It depends on both $\mathbf{z}_t$ and $\mathbf{Z}_{t+1}$ and includes a risk premium that reflects the investors' tolerance to the risky future cash flows, which may be correlated with systemic market-wide risks. More generally for $t \leq w$, we define $\beta_{t,w}(\mathbf{z}_t, \mathbf{Z}_{t+1:w+1})$ to be the discount factor that converts cash flows at the end of period w to equivalent flows at the beginning of period t,

$$\beta_{t,w}(\mathbf{z}_t, \mathbf{Z}_{t+1:w+1}) = \beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) \prod_{j=t+1}^{w} \beta_j(\mathbf{Z}_j, \mathbf{Z}_{j+1}).$$

Stochastic Cost Parameters Since ordering costs (both fixed and variable) are incurred at the beginning of each period, we model them as functions of the realized state-of-the-world information, $\mathbf{z}_t$. We assume that salvage is only possible at the end of the planning horizon, so that the salvage value is a random variable determined by the end-of-horizon state, $\mathbf{Z}_{T+1}$. Note that the fixed and variable purchasing costs may well vary over time due to fluctuating environment (*e.g.*, exchange rate change). Thus, we define

$$\begin{aligned} K_t(\mathbf{z}_t) \quad &= \quad \text{fixed ordering (and/or setup) cost at the beginning of period } t, \\ c_t(\mathbf{z}_t) \quad &= \quad \text{variable purchasing (and/or production) cost at the beginning of period } t, \\ v(\mathbf{Z}_{T+1}) \quad &= \quad \text{unit salvage value for on-hand inventory at the end of the planning horizon.} \end{aligned}$$

Since revenues, holding costs, and backorder costs may be incurred throughout each period, we model them as functions of both $\mathbf{z}_t$ and $\mathbf{Z}_{t+1}$. Likewise, the selling prices can vary, e.g., with the exchange rates or weather conditions, and we denote

$$\begin{aligned} h_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) \quad &= \quad \text{holding cost for carrying one unit of inventory from period } t \text{ to } t+1, \\ &\qquad \text{valued at the end of period } t, \\ \pi_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) \quad &= \quad \text{penalty cost (including the goodwill cost) for one unit of backorders} \\ &\qquad \text{or lost sales at the end of period } t\text{, valued at the end of period } t, \\ p_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) \quad &= \quad \text{selling price in period } t\text{, valued at the end of period } t. \end{aligned}$$

To simplify the presentation, we will often make the dependencies on the state of the world vectors implicit. That is, we will often write K_t, c_t, h_t, p_t, π_t, D_t, β_t, $\beta_{t,w}$ and v. Note that $\beta_{t:w} = \beta_t \cdot \beta_{t+1:w}$. We also define, for general A and B, $A^+ = \max\{0, A\}$, $A^- = (-A)^+$, $A \wedge B = \min\{A, B\} = B - (A - B)^-$, $A \vee B = \max\{A, B\}$, and

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

2.2 Model Formulation

Let $G_t(x_t, y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+1}, \mathbf{D}_{t:t+L})$ denote the cost incurred in period $t+L$, given the realized state (x_t, $\mathbf{z}_t$), the order-up-to inventory level y_t ($\geq x_t$), and the random evolution of state-of-the-world and demand information between periods. Then, for $1 \leq t \leq T-L$,

$$G_t(x_t, y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+1}, \mathbf{D}_{t:t+L})$$

$$\begin{aligned}
&= h_{t+L}\cdot(y_t-D_{t,t+L})^+ + \pi_{t+L}\cdot(y_t-D_{t,t+L})^- \\
&\quad -p_{t+L}\cdot\begin{cases} y_t\wedge D_{t,t+L}, & \text{if LS,}\\ (D_{t+L}+(y_{t-1}-D_{t-1,t+L-1})^- -(y_t-D_{t,t+L})^-), & \text{if PTD,}\\ D_{t+L}, & \text{if PTO,}\end{cases}\\
&= h_{t+L}\cdot(y_t-D_{t,t+L})^+ - p_{t+L}\cdot D_{t+L}\\
&\quad +\begin{cases} (\pi_{t+L}+p_{t+L})(y_t-D_{t,t+L})^-, & \text{if LS,}\\ [(\pi_{t+L}+p_{t+L})(y_t-D_{t,t+L})^- - p_{t+L}(y_{t-1}-D_{t-1,t+L-1})^-], & \text{if PTD,}\\ \pi_{t+L}(y_t-D_{t,t+L})^-, & \text{if PTO.}\end{cases}
\end{aligned}$$

In deriving these expressions, we applied the equality $y_t \wedge D_{t,t+L} = D_{t,t+L} - (y_t - D_{t,t+L})^-$, and made use of our assumption that $L = 0$ under lost sales. The number of units delivered to customers under PTD in period $t + L$ (*i.e.*, the multiplicand of p_{t+L} under PTD in the first equation of the expression above for G_t) merits some discussion. The maximum possible sales in period $t + L$ is equal to the demand in that period plus the backlog from the previous period, $D_{t+L}+(y_{t-1}-D_{t-1,t+L-1})^-$. The actual sales will be this maximum sales minus the ending backlog, $(y_t - D_{t,t+L})^-$. In other words, sales in period $t + L$ will be equal to demand minus the change in backlog, $D_{t+L} + (y_{t-1} - D_{t-1,t+L-1})^- - (y_t - D_{t,t+L})^-$.

Letting $\beta_{T+1} = 1$ and defining

$$\hat{\pi}_t \doteq \pi_t + \begin{cases} p_t, & \text{if LS,}\\ p_t-\beta_{t+1}p_{t+1} & \text{if PTD,}\\ 0, & \text{if PTO,}\end{cases} \tag{1}$$

allow us to rewrite this problem in terms of the modified penalty costs $\hat{\pi}_t$ and a single period cost in period $t + L$, $1 \le t \le T - L$, of

$$\widehat{G}_t(y_t,\mathbf{z}_t,\mathbf{Z}_{t+1:t+L+2},\mathbf{D}_{t:t+L}) = h_{t+L}\ (y_t-D_{t,t+L})^+ + \hat{\pi}_{t+L}\ (y_t-D_{t,t+L})^- - p_{t+L}\ D_{t+L}. \tag{2}$$

Note that the modified penalty cost also captures changes in revenues resulting from shortages. In the PTD case where the modified penalty cost $\hat{\pi}_t$ is a function of both $p_t(\mathbf{z}_t, \mathbf{Z}_{t+1})$ and $p_{t+1}(\mathbf{Z}_{t+1:t+2})$ (p_{T+1} and the end of horizon costs will be defined later), it may be negative for some realizations of the state-of-the-world vectors. This implies that by backlogging instead of satisfying the demand, the firm may reap a higher revenue. The dynamic programming formulation for the expected discounted cost at the beginning of period t, $1 \le t \le T - L$, is given by

$$\begin{aligned}
f_t(x_t,\mathbf{z}_t) &= -c_t\,x_t + \min_{y_t\ge x_t}\{K_t\delta(y_t-x_t)+V_t(y_t,\mathbf{z}_t)\}, && (3)\\
V_t(y_t,\mathbf{z}_t) &= c_t\,y_t + E_{\mathbf{D}_{t:t+L},\mathbf{Z}_{t+1:t+L+2}}[\beta_{t,t+L}\widehat{G}_t(y_t,\mathbf{z}_t,\mathbf{Z}_{t+1:t+L+2},\mathbf{D}_{t:t+L})\\
&\quad +\beta_t f_{t+1}(\gamma(y_t-D_t),\mathbf{Z}_{t+1})], && (4)
\end{aligned}$$

where

$$\gamma(x) \doteq \begin{cases} x^+ & \text{under LS}\\ x & \text{under PTO and PTD}\end{cases}.$$

To calculate the end-of-horizon cost, we assume that all outstanding backorders are filled via a zero-lead-time purchase made at the end of period T at the unit cost of $c_{T-L+1}(\mathbf{Z}_{T+1})$. (The time index $T-L+1$ is chosen for notational convenience.) In the PTD case, the end-of-horizon revenue is received at the beginning of period $T+1$ at $p_{T+1}(\mathbf{Z}_{T+1})$ per unit, and is already included in $\widehat{G}_{T-L}$. Therefore, the end-of-horizon cost that initiates the recursion in (3) and (4) is defined as:

$$
\begin{aligned}
f_{T-L+1}(x_{T-L+1}, \mathbf{z}_{T-L+1}) \quad = \quad & E_{D_{T-L+1,T}, \mathbf{Z}_{T-L+2:T+1}} \Big\{ \beta_{T-L+1,T} \Big[-v \cdot (x_{T-L+1} - D_{T-L+1,T})^{+} \\
& + \begin{cases} 0, & \text{if lost sales,} \\ c_{T-L+1} \cdot (x_{T-L+1} - D_{T-L+1,T})^{-}, & \text{if backorder} \end{cases} \Big\}. \qquad (5)
\end{aligned}
$$

3 Structural Analyses and Results

When all cost parameters and discount factors are known in advance, it is well known that an (s, S) policy is optimal for the above problem, provided that a few side conditions are satisfied. For example, one precondition for the optimality of (s, S) policies is the monotonicity of the fixed ordering cost, $K_t \geq \beta_t \cdot K_{t+1}$, which is necessary for Scarf's (1959) K_t-convexity property to hold. This condition is one of several that may fail when the business environment is particularly volatile.

In this section, we will provide in Theorem 1 sufficient conditions for an (s, S) policy to be optimal, and discuss the implications of these conditions. This will help decision-makers understand when an optimal policy may be non-trivial under multiple uncertainties. We begin this section by discussing some important properties of the functions defined in the previous section. Throughout the paper, all the proofs can be found in the Appendix.

3.1 Properties

We start with the finiteness and continuity of the functions $\widehat{G}_t$, f_t and V_t, and then discuss the asymptotic properties of functions with linear growth rates.

3.1.1 Finiteness and Continuity

So far we have allowed the functions defined in Section 2.2 to take on values in $\Re \cup \{\pm\infty\}$. We now consider the finiteness and continuity of the functions defined above. Let $u_t = E\{\beta_{t,t+L} \cdot [h_{t+L} \vee \beta_{t+L+1} \cdot p_{t+L+1} \vee (\pi_{t+L} + p_{t+L})]\}$ for $1 \leq t \leq T-L$, and let $U_{T-L+1} = E\{\beta_{T-L+1,T}[|v| \vee c_{T-L+1}]\}$. The following finiteness conditions are assumed to hold throughout this paper.

Assumption 1 *$E[\beta_{t,r} \cdot u_r]$, $E\{\beta_{t,w} \cdot D_w \cdot E[\beta_{w+1,r} \cdot u_r]\}$, $E[\beta_{t,T-L} \cdot U_{T-L+1}]$, and $E\{\beta \cdot D_w \cdot E[\beta_{w+1,T-L} \cdot U_{T-L+1}]\}$ are finite for all $\mathbf{z_t}$ and all $1 \leq t \leq w+1$, $1 \leq w \leq r \leq T-L$.*

Assumption 2 *There exists a policy ω with a finite expected cost. Define f_t^ω and V_t^ω to be functions analogous to f_t and V_t in (3) and (4), under policy y_t^ω. (To be specific, we assume that f_t^ω and V_t^ω satisfy (4), and satisfy (3) when we replace the minimum over y_t with the value of y_t^ω under policy ω.) Then $f_t^\omega(x_t, \mathbf{z}_t)$ and $V_t^\omega(y_t^\omega, \mathbf{z}_t)$ are finite for all finite values x_t and y_t^ω, for all $\mathbf{z}_t$.*

Note that Assumption 2 follows necessarily from Assumption 1 only when the discount factors never exceed 1; *i.e.*, only when $\beta_w \leq 1$ for all $w \geq t$ and for all $(t, \mathbf{z_t})$. These finiteness assumptions are needed to establish the finiteness and Lipschitz continuity of the cost functions in Lemma 1. Note that we consider a function to be also Lipschitz-continuous if it is equal to $-\infty$ everywhere.

Lemma 1 *For every time period t and every information set $\mathbf{z}_t$ the following properties hold.*

- *The function $E_{\mathbf{D}_{t:t+L}, \mathbf{Z}_{t+1:t+L+2}}[\beta_{t,t+L}\widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ is a well-defined Lipschitz continuous function of y_t. It is finite if y_t is finite.*
- *$f_t(x_t, \mathbf{z}_t)$ is a Lipschitz continuous function in x_t. If x_t is finite, then f_t is either finite or equal to $-\infty$; i.e., $f_t \in \Re \cup \{-\infty\}$.*
- *$V_t(y_t, \mathbf{z}_t)$ is a Lipschitz continuous function in y_t. If y_t is finite, then $V_t \in \Re \cup \{-\infty\}$.*
- *If the minimum of $V_w(y_w, z_w)$ is attained at a finite y_w for all $w \geq t$ and all $\mathbf{z}_w$, then V_t and f_t are finite when x_t, y_t are finite.*

3.1.2 Asymptotic Properties of Functions with Linear Growth Rates

As one will see, the functions we will deal with are all asymptotically linear. Thus, we define some basic properties of such functions.

Definition 1 *We say that*

1. *$f(x)$ is $[a, \cdot]$-divergent if $\lim_{x \to -\infty}[f(x) - ax] = \infty$;*
2. *$f(x)$ is $[\cdot, b]$-divergent if $\lim_{x \to \infty}[f(x) - bx] = \infty$; and*
3. *$f(x)$ is $[a, b]$-divergent if it is both $[a, \cdot]$-divergent and $[\cdot, b]$-divergent.*

Definition 2 *We say that*

1. *$f(x)$ is $[a, \cdot]$-asymptotic if for all $\epsilon > 0$, $f(x)$ is $[a + \epsilon, \cdot]$-divergent and $-f(x)$ is $[-a + \epsilon, \cdot]$-divergent;*

2. *$f(x)$ is $[\cdot, b]$-asymptotic if for all $\epsilon > 0$, $f(x)$ is $[\cdot, b-\epsilon]$-divergent and $-f(x)$ is $[\cdot, -b-\epsilon]$-divergent;*

3. *$f(x)$ is $[a, b]$-asymptotic if it is $[a, \cdot]$-asymptotic and $[\cdot, b]$-asymptotic.*

The definitions of $[a, b]$-divergent and $[a, b]$-asymptotic will be used in the proofs of Lemmas 3 and 4. The concept of $[a, b]$-asymptoticness describes how close a function is to a linear function and provides an easy way to identify whether a function is [0,0]-divergent. For example, if a function is $[a, \cdot]$-asymptotic and $a < 0$, then, by Definition 2(1), we know that it must be $[0, \cdot]$-divergent. Lemma 3 will show that [0,0]-divregence of the cost function is essential for guaranteeing a finite reorder point s and a finite order-up-to point S. For example, if a cost function $f(x)$ is $[0, \cdot]$-divergent, then the value of the function will go to infinity when x goes to $-\infty$, which guarantees a finite s. Below we provide some properties of $[a, b]$-divergent and $[a, b]$-asymptotic functions.

Property 1 1. *If $f_i(x)$ is $[a_i, b_i]$-divergent and $c_i \geq 0$, then $\sum_i c_i f_i(x)$ is $\left[\sum_i c_i a_i, \sum_i c_i b_i\right]$-divergent.*

2. *If $f_i(x)$ is $[a_i, b_i]$-asymptotic, then $\sum_i c_i f_i(x)$ is $\left[\sum_i c_i a_i, \sum_i c_i b_i\right]$-asymptotic.*

Lemma 2 *Let $f(x)$ be $[a, b]$-asymptotic, and let W and D be correlated random variables with $E(|W|) > 0$ and either $W \geq 0$ or $W \leq 0$. Suppose that $E(W)$, $E(WD)$, and $E[Wf(x-D)]$ exist. Then, $E[Wf(x-D)]$ is $[E(W)a, E(W)b]$-asymptotic if $E[WD]$ and $E[Wf(x-D)]$ are finite for all finite x.*

Both Property 1 and Lemma 2 will be used later to prove the asymptoticness of the cost functions. For example, one consequence of Lemma 2 is that $E_{\mathbf{D}_{t:t+L}, \mathbf{Z}_{t+1:t+L+2}}[\beta_{t,t+L}\widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ is $[E_{\mathbf{Z}_{t+1:t+L+2}}(-\hat{\pi}_{t+L}), E_{\mathbf{Z}_{t+1:t+L+1}}(h_{t+L})]$-asymptotic.

3.2 Sufficient Properties for the Optimality of (s, S) Policies

Let

$$S_t(\mathbf{z}_t) = \arg\min\{V_t(y_t, \mathbf{z}_t)\}, \tag{6}$$

$$s_t(\mathbf{z}_t) = \inf\{y_t : K_t(\mathbf{z}_t) + V_t(S_t(\mathbf{z}_t), \mathbf{z}_t) \geq V_t(y_t, \mathbf{z}_t)\}. \tag{7}$$

Since $S_t(\mathbf{z}_t)$ satisfies the inequality in (7), $s_t(\mathbf{z}_t)$ is well-defined and $S_t(\mathbf{z}_t) \geq s_t(\mathbf{z}_t)$. Also note that under case LS, $s_t(\mathbf{z}_t) \geq 0$. Although the domain of $V_t(\cdot, \mathbf{z}_t)$ is limited to real numbers, we relax this restriction to allow both $S_t(\mathbf{z}_t)$ and $s_t(\mathbf{z}_t)$ to be infinite. In the event that $s_t(\mathbf{z}_t) = -\infty$, no order is ever placed in period t (which may indeed happen in practice), and which we do not consider to be pathological. On the other hand, if $S_t(\mathbf{z}_t) = \infty$, then all orders placed are for infinite amounts of inventory and an infinite profit may be possible; we consider such problems to be poorly-posed. In this paper, the phrase *an (s, S) policy is optimal* implies that $S_t(\mathbf{z}_t) < \infty$ in (6). To simplify the presentation, we will often make the time index and the dependence on the state of the world implicit; that is, we will often write (s, S) instead of $(s_t(\mathbf{z}_t), S_t(\mathbf{z}_t))$.

The following lemma provides the sufficient conditions for the optimality of an (s, S) policy.

Lemma 3 *Suppose that $V_t(\cdot, \mathbf{z}_t)$ is K_t-convex and $[\cdot, 0]$-divergent. Then in period t with information set $\mathbf{z}_t$, an (s, S) policy will be optimal. In cases PTD and PTO, if in addition $V_t(\cdot, \mathbf{z}_t)$ is $[0, \cdot]$-divergent, then $s > -\infty$ under the optimal policy.*

While K_t-convexity of the cost function guarantees the optimality of an (s, S) policy, $[0, \cdot]$-divergence ($[\cdot, 0]$-divergence) determines whether s (S) is finite. In traditional inventory models with constant cost parameters and discount factors, $[0, 0]$-divergence of cost functions holds naturally. In our generalization, it is possible that $[0, 0]$-divergence is violated, which could lead to infinite policy values.

The optimality of an (s, S) policy is thus reduced to the K_t-convexity and $[a, b]$-divergence of the functions $V_t(\cdot, \mathbf{z}_t)$. We consider these two properties in turn, starting with $[a, b]$-divergence.

3.2.1 $[a, b]$-Divergence of the Function $V_t(\cdot, \mathbf{z}_t)$

Let

$$\mathcal{A}_t = \mathcal{A}_t(\mathbf{z_t}) = \begin{cases} c_t + E_{\mathbf{Z}_{t+1:t+L+2}}[-\beta_{t,t+L} \cdot \hat{\pi}_{t+L} + \beta_t(-c_{t+1} + \mathcal{A}^+_{t+1}(\mathbf{Z}_{t+1}))], & \text{if backorder,} \\ c_t - E_{\mathbf{Z}_{t+1:t+L+2}}[\beta_{t,t+L} \cdot \hat{\pi}_{t+L}], & \text{if lost sales,} \end{cases} \tag{8}$$

$$\mathcal{B}_t = \mathcal{B}_t(\mathbf{z_t}) = c_t + E_{\mathbf{Z}_{t+1:T+1}} \left[\sum_{j=t+L}^{T} \beta_{t,j} \cdot h_j - \beta_{t,T} \cdot v \right], \tag{9}$$

for $1 \le t \le T - L$, with $\mathcal{A}_{T-L+1} = 0$. In the backorder cases (PTD and PTO), $\mathcal{A}_t$ is the expected cost differential, as of the beginning of period t, between the options of (a) ordering one unit in period t at c_t, and (b) carrying one unit of backorder in period $t + L$ at the modified backorder cost $\hat{\pi}_{t+L}$ and then ordering the unit in period $t+1$ at c_{t+1} (or even later if that would be cheaper still; *i.e.*, if $\mathcal{A}_{t+1} > 0$). Thus we expect $A_t < 0$. In case LS, $\mathcal{A}_t$, for $t < T$, is irrelevant because the domain of $V_t(\cdot, \mathbf{z}_t)$ is $\Re^+$. Intuitively, $\mathcal{B}_t$ is the expected cost of ordering one unit in period t, carrying it from period $t + L$ to the end of the horizon, and salvaging it. If $\mathcal{B}_t < 0$ for some t

and $\mathbf{z}_t$, then an infinite amount of inventory should be ordered, resulting in an expected cost of $-\infty$. To preclude such ill-posed problems, we require that $\mathcal{B}_t > 0$ for all realizations $\mathbf{z}_t$ of $\mathbf{Z}_t$, for $1 \leq t \leq T - L$.

Lemma 4 *$V_t(\cdot, \mathbf{z}_t)$ is $[\mathcal{A}_t, \mathcal{B}_t]$-asymptotic in cases PTD and PTO, and $[\cdot, \mathcal{B}_t]$-asymptotic in case LS.*

The asymptoticness of the cost functions is directly related to the finiteness of s and S. By Definitions 1 and 2, $V_t(\cdot, \mathbf{z}_t)$ is $[0, \cdot]$-divergent and a finite reorder point is guaranteed if $\mathcal{A}_t < 0$. On the other hand, if $\mathcal{B}_t > 0$, then $V_t(\cdot, \mathbf{z}_t)$ is $[\cdot, 0]$-divergent and S is finite.

3.2.2 The K_t-Convexity of $V_t(\cdot, \mathbf{z}_t)$

We now consider the property required for the functions $V_t(\cdot, \mathbf{z}_t)$ to be K_t-convex.

Lemma 5 *$V_t(y_t, \mathbf{z}_t)$ is K_t-convex if the following conditions hold. For all possible $\mathbf{z}_{T-L}$, $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$. In addition, for all $\mathbf{z}_j$ and all $j = t, \cdots, T-L-1$,*

1. *$h_{j+L} + \hat{\pi}_{j+L} \geq 0$ with backorders and $h_j + \hat{\pi}_j - c_{j+1} \geq 0$ with lost sales, and*
2. *$K_j(\mathbf{z}_j) \geq E_{\mathbf{Z}_{j+1}}[\beta_j(\mathbf{z}_j, \mathbf{Z}_{j+1}) K_{j+1}(\mathbf{Z}_{j+1})]$.*

We will discuss the implications of the conditions in Lemma 5 shortly after we introduce Theorem 1 in the following section.

3.3 Sufficient Conditions for the Optimality of (s, S) Policies

Combining Lemmas 3, 4 and 5, we summarize the conditions under which (s, S) policies are optimal in the following theorem.

Theorem 1 *An (s, S) policy is optimal in period t, if*

1. *$\mathcal{B}_j > 0$ for all realizations $\mathbf{z}_j$ and all $j = t, \cdots, T-L$,*
2. *$\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$ for all possible realizations $\mathbf{z}_{T-L}$, and*
3. *for all realizations $\mathbf{z}_j$ and all $j = t, \cdots, T-L-1$,*
   (a) *$h_{j+L} + \hat{\pi}_{j+L} \geq 0$ under backlogging and $h_j + \hat{\pi}_j - c_{j+1} \geq 0$ under lost sales; and*
   (b) *$K_j(\mathbf{z}_j) \geq E_{\mathbf{Z}_{j+1}}[\beta_j(\mathbf{z}_j, \mathbf{Z}_{j+1}) \cdot K_{j+1}(\mathbf{Z}_{j+1})]$.*

If in addition $\mathcal{A}_t < 0$, then $s > -\infty$.

With the exception of condition (3b), the conditions in Theorem 1 are all related to the effects of uncertainties on the holding and backorder costs. While Conditions 1 and 2 can hold with appropriately chosen salvage values, the rest of them can fail in practice. The unlabeled condition $\mathcal{A}_t < 0$ is not required for the optimality of an (s, S) policy but ensures that a positive amount of inventory will be ordered if the inventory level is low enough. We turn now to discuss the implications of each condition in the Theorem in further detail.

Condition 1, $\mathcal{B}_t > 0$: This condition will generally hold for reasonable end-of-horizon salvage values. If $\mathcal{B}_t \le 0$, it would be optimal to order an infinite amount of inventory in period t, hold it, and salvage it at the end of the horizon; in that case, the problem is ill-posed, probably because the salvage value v was overestimated.

To the extent that the salvage value is chosen to approximate the value of inventory beyond the given finite horizon, we suggest that an appropriate stochastic value for v that minimizes the end-of-horizon effect should satisfy

$$v(\mathbf{z}_{1:T+1}) < \min\left\{\frac{1}{\beta_{t,T}(\mathbf{z}_{t:T+1})}\left[c_t(\mathbf{z}_t) + \sum_{j=t+L}^{T}\beta_{t,j}(\mathbf{z}_{t:j+1})\,h_j(\mathbf{z}_{j:j+1})\right] : 1 \le t \le T-L\right\} \tag{10}$$

for all $\mathbf{z}_{1:T+1}$, when viewed from the beginning of period $T+1$. (10) ensures that $\mathcal{B}_t > 0$. Note that v may well depend on the entire historical state of the world $(\mathbf{z}_1, \mathbf{z}_2, ..., \mathbf{z}_{T+1})$ instead of just $\mathbf{z}_{T+1}$ as we have formally assumed. However, we can assume without loss of generality that $\mathbf{z}_{T+1}$ contains all of the information $(\mathbf{z}_1, \mathbf{z}_2, ..., \mathbf{z}_{T+1})$, so all of our results continue to hold.

Condition 2, $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \ge 0$: This condition guarantees the K_t-convexity of the cost function in the last period of the horizon and so the optimality of an (s, S) policy. Although this condition generally holds, it can fail for inappropriate end-of-horizon parameters are. In that case, the modeler should restrict the applicability of our model to appropriate end-of-horizon parameters that ensure this condition to hold. More specifically, since $L = 0$ in the lost sales case, $\mathcal{A}_T = c_T - E_{\mathbf{Z}_{T+1}}[\beta_T(p_T + \pi_T)] \le 0$ if it is cheaper to purchase a unit in period T and deliver it to a client than it is to lose the sale. If p_T and π_T are observed at the beginning of period T, then $\mathcal{A}_T$ is naturally negative and we have $\mathcal{B}_T - \mathcal{A}_T \ge 0$. However, in the real world, p_T and π_T represent cash flows that are realized at various times within period T. Consequently, p_T and π_T may be dependent on $\mathbf{Z}_{T+1}$ as well as $\mathbf{z}_T$, in which case $\mathcal{A}_T > 0$ can conceivably occur. In that case, $\mathcal{B}_T - \mathcal{A}_T \ge 0$ if it is appropriate to specify the salvage value of v such that

$$v(\mathbf{z}_{T:T+1}) \le h_T(\mathbf{z}_{T:T+1}) + \hat{\pi}_T(\mathbf{z}_{T:T+1})$$

and (10) holds as well.

Under backlogging, $\mathcal{A}_{T-L} = c_{T-L} - E_{\mathbf{Z}_{T-L+1:T+1}}[\beta_{T-L,T} \cdot (c_{T-L+1} + \hat{\pi}_T)] \leq 0$ is equivalent to specifying that it is cheaper to purchase a unit in period $T - L$ and deliver it to a client in period T than it is to buy and deliver the unit at the end of period T at cost of $c_{T-L+1} + \pi_T$. (Recall that c_{T-L+1} is the unit cost of a unique, zero-lead-time purchase opportunity that is available only at the end of period T.) Only certain values of c_{T-L+1} will ensure $E(\beta_{T-L,T} \cdot c_{T-L+1}) \geq c_{T-L} - E(\beta_{T-L,T} \cdot \hat{\pi}_T)$, so as to reduce the end-of-horizon effect. We recommend defining c_{T-L+1} so that

$$c_{T-L+1}(\mathbf{z}_{T-L:T+1}) \geq c_{T-L}(\mathbf{z}_{T-L})/\beta_{T-L+1,T}(\mathbf{z}_{T-L:T+1}) - \hat{\pi}_T(\mathbf{z}_T, \mathbf{z}_{T+1}).$$

This potentially makes c_{T-L+1} a function of $\mathbf{z}_{T-L:T+1}$ rather than a function of $\mathbf{z}_{T+1}$ as we originally assumed, but as we mentioned before, we can assume without loss of generality that the information in $\mathbf{z}_{T-L:T+1}$ is included in $\mathbf{z}_{T+1}$.

Condition (3a) , $h_t + \hat{\pi}_t \geq 0$ (backlogging) and $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ (lost sales): This condition guarantees the convexity of the (transformed if LS) single-period inventory related cost function, $h_t(y_t - D_t)^+ + \hat{\pi}_t(y_t - D_t)^-$ under backlogging and $(h_t - c_{t+1})(y_t - D_t)^+ + \hat{\pi}_t(y_t - D_t)^-$ under lost sales, which is critical for the K_t-convexity of the cost functions. This condition corresponds to the situation in which we have a unit of demand in period t, and the inventory to meet that demand. However, under backlogging, we may sometimes deliberately choose to carry the inventory into period $t+1$ and deliver it to the client then. In the lost sales case, we may choose to lose customer demand and use the unit of inventory to displace a unit that we would otherwise have purchased in period $t+1$. This condition states that such decisions are not preferable. The condition always holds for PTO systems ($\hat{\pi}_t = \pi_t$), and is a very sensible assumption in most circumstances. However, this condition can fail, especially for PTD systems.

Condition (3b) , $K_t(\mathbf{z}_t) \geq E_{\mathbf{Z}_{t+1}}[\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) \cdot K_{t+1}(\mathbf{Z}_{t+1})]$: This condition clearly holds in the absence of fixed costs; *i.e.*, when $K_t = 0$ for all t. When fixed costs are present, uncertainties can cause it to fail, in which case (s, S) policies may not be optimal.

Condition $\mathcal{A}_t < 0$: While this condition is not needed for an (s, S) policy to be optimal, it guarantees that $s > -\infty$; *i.e.*, that in period t there exists an inventory level below which an order would be placed. Under lost sales, the comments in the first paragraph of our discussion of Condition 2 regarding $\mathcal{A}_T < 0$ apply. The condition generally holds, although (depending on the specific model assumptions), uncertanties might cause it to fail.

Under backlogging, this condition states that uncertainties are not high enough to motivate delaying ordering and order fulfillment. If $\mathcal{A}_t \geq 0$, then for some time period $w \geq t$, $\mathcal{A}_{w+1} < 0 < \mathcal{A}_w$. This implies that $c_w - E[\beta_{w,w+L} \cdot \pi_{w+L} + \beta_w \cdot c_{w+1}] \geq 0$. In other words, the expected cost of delaying a purchase to period $w+1$ and delivering it to the client one period late (in period $w+L+1$) is lower than the expected cost of making the purchase in period w and delivering it to the client on time in period $w+L$. Although $\mathcal{A}_t < 0$ for all t is a sensible and very common assumption, uncertainties may on occasion cause it to fail.

We have been consistent in asserting that a condition holds only if it holds for all realizations of the state of the world. Thus, a firm will include various uncertainties in costs even when the state of the world is stationary, although doing so is more important under non-stationary environment. The impact of the failure of the sufficient conditions is two-fold. First, an (s, S) policy may no longer be optimal, in which case non trivial optimal policies can only be obtained by solving a dynamic programming problem in each period. Second, even when an (s, S) or order-up-to policy remains optimal, the optimal order-up-to levels can be sharply different from those under stable environment. For example, if Condition $\mathcal{A}_t < 0$ fails, it may be optimal to never order.

4 What Happens When Sufficient Conditions Fail

In this section we first provide analytical examples, in Sections 4.1 and 4.2, of non-(s, S) policies that are optimal when the sufficient conditions in Section 3.3 fail to hold. We then demonstrate in Section 4.3 that (s, S) policies can fail to be optimal in randomly-generated problems when the sufficient conditions in Section 3.3 are violated. Since conditions 1 and 2 of Theorem 1 can be made to hold by choosing appropriate end-of-horizon parameters, and violation of condition $\mathcal{A}_t < 0$ alone generates trivial policies in which orders are never placed, we focus on conditions (3a) and (3b).

4.1 Analytical Examples of Non-Convexity When Condition (3a) Fails

Throughout this subsection we assume that the fixed cost K_t is zero. When Condition (3a) fails, the cost function may not be convex, in which case an order-up-to policy may no longer be optimal. In the dynamic programming recursion (3) – (4), there are only two ways a non-convexity can be introduced. Either $\widehat{G}_t(y)$ can fail to be convex in (4), or the nonlinearity of $\gamma(y - D_t)$ can introduce a non-convexity in the Lost Sales case. For the PTO case, $\gamma(y - D_t) = y - D_t$ so the γ function cannot cause a non-convexity. For $\widehat{G}_t(y) = h_t(y - D)^+ + \pi_t(D_t - y)^+$ to fail to be convex, either $h_t < 0$ or $\pi_t < 0$ would have to occur with positive probability, which is unlikely. As such, we will only consider two cases – Lost Sales and PTD. In both cases, in our examples Condition (3a) holds in expectation, but not with probability 1 as is required. As a result, an order-up-to policy may not be optimal in some periods.

Under Lost Sales

Consider a two period problem in which all parameters are deterministic except for c_1 and p_1. Define $v = 0$, $L = 0$, $D_1 = 2$, $D_2 = 1$, $c_2 = 4$, $p_2 = 5$, and $\beta_t = 1$, $\pi_t = 0$ and $h_t = 1$ for $t \in \{1, 2\}$. We assume that c_1 and p_1 are random and governed by an underlying variable Z where $P(Z = 0) = 2/3$ and $P(Z = 1) = 1/3$. Z is observed at the beginning of period 1. When $Z = 0$ then $(c_1, p_1) = (4, 5)$, while when $Z = 1$ then $(c_1, p_1) = (2, 1)$. Note that $\hat{\pi}_t = p_t$.

We can see that condition (3a) fails in period 1 when $Z = 1$, since $h_1 + \hat{\pi}_1 - c_2 = 1 + 1 - 4 < 0$. However, it holds in expectation because $E[h_1 + \hat{\pi}_1 - c_2] = 1 + (\frac{2}{3} \times 5 + \frac{1}{3} \times 1) - 4 > 0$. It can be shown that when $Z = 1$,

$$V_1(y_1) = \begin{cases} y_1 - 1, & \text{if } y_1 \leq 2, \\ -y_1 + 3, & \text{if } 2 < y_1 \leq 3, \\ 4y_1 - 12, & \text{if } y_1 > 3, \end{cases}$$

and so the optimal ordering strategy in period 1 is

$$y_1^*(x_1) = \begin{cases} x_1, & \text{if } x_1 \leq 1 \text{ or } x_1 \geq 3, \\ 3, & \text{if } 1 < x_1 < 3, \end{cases}$$

which is not an order-up-to policy. If $Z = 0$, the optimal policy in period 1 is order-up-to.

Under PTD

Again, consider a two period problem in which all parameters are deterministic except for (c_1, p_1). We have $v = 0$, $L = 0$, $D_t = \beta_t = h_t = \pi_t = 1$, and for $t \in \{2, 3\}$ we have $p_t = 6$ and $c_t = 3$. As before, c_1 and p_1 are governed by an underlying variable Z where $P(Z = 0) = 2/3$, $P(Z = 1) = 1/3$, and Z is observed at the start of period 1. When $Z = 0$ then $(c_1, p_1) = (3, 6)$, and when $Z = 1$ then $(c_1, p_1) = (1, 2)$.

When $t = 1$, Condition (3a) becomes $0 \leq h_1 + \hat{\pi}_1 = h_1 + \pi_1 + p_1 - p_2 = p_1 - 4$. If $Z = 0$ this condition holds as $p_1 - 4 = 2$, and an order-up-to policy is optimal. When $Z = 1$ this condition fails as $p_1 - 4 = -2 < 0$. In expectation it holds, as $E[p_1 - 4] = (\frac{2}{3} \times 6 + \frac{1}{3} \times 2) - 4 > 0$. We can show that when $Z = 1$,

$$V_1(y_1) = \begin{cases} y_1 - 5, & \text{if } y_1 \leq 1, \\ -y_1 - 3, & \text{if } 1 < y_1 \leq 2, \\ 3y_1 - 11, & \text{if } y_1 > 2, \end{cases}$$

and the optimal ordering strategy in period 1 is

$$y_1^*(x_1) = \begin{cases} x_1, & \text{if } x_1 \le 0 \text{ or } x_1 \ge 2, \\ 2, & \text{if } 0 < x_1 < 2, \end{cases}$$

which is not an order-up-to policy.

In both examples, the cost function $V_1(y_1)$ displays an $up - down - up$ pattern and the order-up-to policy is optimal only if the inventory level is above a threshold. For the PTD example, if the inventory level is below the threshold there are either existing backorders carried over from before period 1 or potential new backorders in the current period. The optimal policy is to not order, and to meet those backorders in the next period at a higher selling price.

4.2 Analytical Examples of Non-K_t-Convexity When Condition (3b) Fails

When $K_t > 0$ and Condition (3b) fails, the cost function may not be K_t-convex and an (s, S) policy may or may not be optimal. Under all three stockout scenarios (PTO, PTD, and LS), we consider the following two period problem where all parameters are deterministic except K_2. We set $v = 0$, $L = 0$, $K_1 = 1$, $c_t = 0$, $h_t = 1$, $\hat{\pi}_t = 2$, $D_t = 4$, and $\beta_t = 1$. We note that for each of the scenarios PTO, PTD, and LS, it is easy to come up with reasonable values of p_t, π_t that are compatible with $\hat{\pi}_t = 2$. For the sole purpose of simplicity[3] we define $p_t = 0$ and $\pi_t = 2$. K_2 is random and governed by an underlying variable Z with $P(Z = 0) = 20/21$ and $P(Z = 1) = 1/21$. We observe Z at the start of period 1. If $Z = 0$ then $K_2 = 0.8$, and $K_2 = 4$ if $Z = 1$.

We note tha Condition (3b) holds in expectation, as $E(K_2) = 4 \times \frac{1}{21} + 0.8 \times \frac{20}{21} = \frac{20}{21} < K_1$. But it fails when $Z = 1$, in which case

$$V_1(y_1) = \begin{cases} -2y_1 + 12, & \text{if } y_1 \le 4, \\ y_1, & \text{if } 4 < y_1 \le 6, \\ -y_1 + 12, & \text{if } 6 < y_1 \le 8, \\ 2y_1 - 12, & \text{if } y_1 > 8. \end{cases}$$

The optimal ordering strategy in period 1 is

$$y_1^*(x_1) = \begin{cases} 4 \text{ or } 8, & \text{if } x_1 \le 3\frac{1}{2}, \\ x_1, & \text{if } 3\frac{1}{2} < x_1 \le 5, \\ 8, & \text{if } 5 < x_1 \le 7, \\ x_1, & \text{if } x_1 > 7. \end{cases}$$

That is, $V_1(y_1)$ is not K_1-convex, and the optimal ordering strategy is not an (s, S) policy. If $Z = 0$ the optimal policy is (s, S).

[3] Otherwise we would have to describe $V_1(y_1)$ for each of the three stockout scenarios separately.

4.3 Numerical Study

The two-period examples in the previous subsection demonstrate that the introduction of some uncertainties into the first period of a stationary deterministic problem can cause the sufficient conditions to fail, so that the optimal policy can become a non-standard one. In this section we consider a more comprehensive example under the PTD protocol with seven periods where some periods are more uncertain than others. This example shows that there can be stochastic settings for which the sufficient conditions can routinely fail, so that (s, S) policies are suboptimal.

4.3.1 Example Description

To provide some context for this example, we consider a Mexican exporter who sells in the U.S. The financial state vector $\mathbf{Z_t}$ is two dimensional and $\{Z_{1t}, Z_{2t}\}$ ={exchange rate (pesos per dollar), discount factor}. Z_{1t} and Z_{2t} are Markov chains and, for the sole purpose of simplicity, we assume they have the same probability transition matrix. The state spaces and transition probabilities in periods 3 to 6 (*turbulent* periods) and periods 1, 2, and 7 (*tranquil* periods) are given in Table 1, where $e = 10.3$ is the base exchange rate. Note that the discount factors are smaller in turbulent periods as a higher risk premium is typically required for discounting more risky future cash flows. We allow the exchange rates in the turbulent periods to be up to twice as high as those in tranquil periods. For example, during the Mexican presidential election in the 1990's, the Mexican peso was devalued more than 50%.

Periods	**1, 2, 7**	**3 - 6**
Exchange Rates	(0.9e, 1.0e, 1.1e)	(1.6e, 1.9e, 2.2e)
Discount Factors	(0.85, 0.9, 0.95)	(0.7, 0.75, 0.8)
Transition Probabilities	$\begin{pmatrix} 0.7 & 0.2 & 0.1 \\ 0.2 & 0.7 & 0.1 \\ 0.1 & 0.2 & 0.7 \end{pmatrix}$	$\begin{pmatrix} 0.4 & 0.4 & 0.2 \\ 0.3 & 0.4 & 0.3 \\ 0.2 & 0.4 & 0.4 \end{pmatrix}$

Table 1: State Spaces and Probability Transition Matrices for Exchange Rates and Discount Factors

The fixed cost and unit purchasing cost are incurred in Mexico (and denominated in pesos) at the beginning of each period. The holding cost, stockout cost, and revenue are all incurred in the U.S. at the end of the period; they are denominated in dollars but are converted into pesos immediately. Hence, the peso value of the dollar-denominated costs and revenues is proportional to the exchange rate Z_1, as is shown in Table 2.

t	K_t	c_t	h_t	π_t	p_t
1, 2, 7	1	1	$0.01Z_{1,t+1}$	$0.04Z_{1,t+1}$	$0.3Z_{1,t+1}$
3-6	2	1.4	$0.01Z_{1,t+1}$	$0.04Z_{1,t+1}$	$0.3Z_{1,t+1}$

Table 2: Cost Parameters

The salvage values $v(\mathbf{Z_8})$ and $c_{T-L+1}(\mathbf{Z_8})$ at the end of period 7 are specified as in Section 3.3. Demand in each period is assumed to be independent of $\mathbf{Z_t}$ and uniformly distributed over {0,1,2,3,4,5}.

4.3.2 How Often the Conditions Fail and (s, S) Policies Fail to be Optimal

We investigate two cases, one with fixed ordering costs and one without. In the latter case, the (s, S) policy reduces to an order-up-to policy, and only condition (3a) is relevant. We find that conditions (3a) and (3b) fail only in period 2. However, they fail for all nine realizations of (Z_1, Z_2). The first row of Table 3 reports the number of realizations in which conditions (3a) and (3b) fail, and the number of realizations where order-up-to/(s, S) policies fail to be optimal in period 2, for the above example .

Exchange Rates	**Counts of the number of realizations (out of 9 possible) when**			
	Condition (3a) fails	**Condition (3b) fails**	$\mathbf{K_t = 0}$ **and order-up-to policy is not optimal**	$\mathbf{K_t > 0}$ **and** $\mathbf{(s, S)}$ **policy is not optimal**
(1.6e, 1.9e, 2.2e)	9	9	9	9
(1.3e, 1.6e,1.9e)	9	9	2	0
(1.0e, 1.3e,1.6e)	9	9	0	0

Table 3: The Number of Realizations where Conditions (3a) and (3b) Fail and the Number of Realizations where order-up-to/(s, S) Policies are not Optimal

As we can see, in all nine realizations of (Z_1, Z_2) order-up-to policies fail to be optimal in the absence of the fixed costs, and (s, S) policies fail to be optimal in their presence. We find that without the fixed costs, the cost functions exhibit an *up-down-up* pattern in all the cases, and the optimal policy is similar to those in Section 4.1. With fixed costs the optimal policy is characterized by two thresholds on the initial inventory (q_1, q_2), and a single order-up-to level Q. If the initial inventory level is below q_1 or above q_2, then no order should be placed. Otherwise order up to Q. This ordering policy is similar to, but simpler than, the optimal policy in Section 4.2.

4.3.3 Sensitivity Analysis

In this section we first test two new combinations of the exchange rates $\{(1.3e, 1.6e, 1.9e), (1.0e, 1.3e, 1.6e)\}$ for the turbulent periods. The results are shown in lines 2 and 3 of Table 3. We can see that as the exchange rate decreases from $(1.6e, 1.9e, 2.2e)$ to $(1.3e, 1.6e, 1.9e)$, although conditions (3a) and (3b) still fail 100% of the time, order-up-to/(s, S) policies are much more likely to be optimal. This suggests that the degree of the non-stationarity has a large impact on the necessity of the sufficient conditions in guaranteeing the optimality of order-up-to/(s, S) policies, which agrees with our intuition.

We then test different probability transition matrices as given in Table 4, which represent different levels of correlation between $\mathbf{Z_t}$ and $\mathbf{Z_{t+1}}$ for the turbulent periods. In the original example, the correlation between $\mathbf{Z_t}$ and $\mathbf{Z_{t+1}}$ is 0.2. We find that the results change little from the original example except when the correlation coefficient is equal to 0.8 or 1 and there are no fixed costs. At these two correlation levels, only eight of the nine realizations of (Z_1, Z_2) exhibit an up-down-up cost function when condition (3a) fails in period 2. In the ninth realization, the cost function is increasing so that the optimal policy is to never order. This seems to suggest that the correlation between $\mathbf{Z_t}$ and $\mathbf{Z_{t+1}}$ has little impact on the sufficiency/necessity of the conditions.

Correlation Coefficient	0.4	0.6	0.8	1
Transition Probabilities	$\begin{pmatrix} 0.55 & 0.3 & 0.15 \\ 0.225 & 0.55 & 0.225 \\ 0.15 & 0.3 & 0.55 \end{pmatrix}$	$\begin{pmatrix} 0.7 & 0.2 & 0.1 \\ 0.15 & 0.7 & 0.15 \\ 0.1 & 0.2 & 0.7 \end{pmatrix}$	$\begin{pmatrix} 0.85 & 0.1 & 0.05 \\ 0.075 & 0.85 & 0.075 \\ 0.05 & 0.1 & 0.85 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$

Table 4: Probability Transition Matrices in Turbulent Periods with Different Correlation Levels

5 Conclusion

In this paper we generalize the traditional inventory models to allow the discount factors, operational costs, and demand to be stochastic and evolve according to the state of the world. We prove that for all three stockout protocols (PTO, PTO, and LS), a state-dependent (s, S) ordering policy will be optimal under certain circumstances. In stochastic business environment, when an (s, S) policy is optimal, neither s nor S is necessarily finite. We also provide the conditions under which both s and S will be finite. The conditions that guarantee the optimality and the finiteness of the (s, S) policy are then discussed to help managers understand the underlying trade-off for their inventory decision making. Furthermore, we provide both analytical and numerical examples of non-trivial policies when the sufficient conditions fail to hold. Our numerical study suggests that if the system's degree of non-stationarity is high, then the sufficient conditions are also necessary to guarantee the optimality of the (s, S) policies. Otherwise, even if the sufficient conditions fail, the (s, S) policies can still be optimal. Furthermore, the correlation in the financial state over time does not seem to affect the sufficiency and necessity of the conditions.

References

[1] Cheng, F. and S. P. Sethi. 1999. Optimality of State-Dependent (s, S) Policies in Inventory Models with Markov-Modulated Demand and Lost Sales. *Production and Operations Management*, 8(2),183-192.

[2] Fabian, T., J. L. Fisher, M. W. Sasieni, and A. Yardeni. 1959. Purchasing Raw Material on a Fluctuating Market. *Operations Research*, 7(1),107-122.

[3] Gaur, V., and S. Seshadri. 2005. Hedging Inventory Risk through Market Instruments. *Manufacturing & Service Operations Management*, 7(2), 103-120.

[4] Graves, S. 1999. A Single-Item Inventory Model for a Nonstationary Demand Process. *Manufacturing and Service Operations Management*. 1(1), 50-61.

[5] Iglehart, D., and S. Karlin. 1962. Optimal Policy for Dynamic Inventory Process with Nonstationary Stochastic Demands. *Studies in Applied Probability and Management Science*. Edited by K. Arrow, S. Karlin, and H. Scarf. Stanford University Press, Stanford, California.

[6] Johnson G. D., and H. E. Thompson. 1975. Optimality of Myopic Inventory Policies for Certain Dependent Demand Processes. *Management Science*, 21(11), 1303-1307.

[7] Kalynon, B. A. 1971. Stochastic Prices in a Single-Item Inventory Purchasing Model. *Operations Research*, 19(6), 1434-1458.

[8] Karlin, S. and A. Febens. 1959. The (s,S) Inventory Model under Markovian Demand Process. *Mathematical Methods in the Social Sciences*. Edited by J. Arrow, S. Karlin, and P.Suppes, Stanford University Press, Stanford, CA.

[9] Scarf, H. E. 1959. The Optimality of (S,s) Policies in the Dynamic Inventory Problem. Chapter 13, in *Mathematical Methods in the Social Sciences*, ed. by K.J. Arrow, S. Karlin and P. Suppes. Stanford University Press.

[10] Sethi S. P. and F. Cheng. 1997. Optimality of (s, S) Policies in Inventory Models with Markovian Demand. *Operations Research*, 45(6), 931-939.

[11] Song, J., and P. Zipkin. 1993. Inventory Control in a Fluctuating Demand Environment. *Operations Research*, 41(2), 351-370.

[12] Veinott, A., Jr. 1966. On the Optimality of (s, S) Inventory Policies: New Conditions and a New Proof. *SIAM Journal on Applied Mathematics*, 14(5), 1067-1083.

Appendix to "New Sufficient Conditions for $(s,\ S)$ Policies to be Optimal in Systems with Multiple Uncertainties"

Proof of Lemma 1: We prove a series of claims, the first of which has to do with the quarter-t cost function $\widehat{G}_t$. Let $1 \le \tau \le t$. The Finiteness Assumption implies that for every $\mathbf{z}_\tau$, $E_{\mathbf{D}_{t:t+L},\mathbf{Z}_{\tau+1:t+L+2}}[\beta_{\tau,t+L}\widehat{G}_t(y_t,\mathbf{z}_t,\mathbf{Z}_{t+1:t+L+2},\mathbf{D}_{t:t+L})]$ exists and is finite when $y_t = 0$. Equations (1) and (2) imply that the derivative of $\widehat{G}_t$, before taking any expectations, is between $-(\pi_t+p_t)$ and $h_t \vee \beta_{t+1}\, p_{t+1}$. By the Finiteness Assumption, if either $\tau = t$ and y_t is finite, or if $y_t = (x_1 - D_{1,t-1})^+$, then $E_{\mathbf{D}_{t:t+L},\mathbf{Z}_{\tau+1:t+L+2}}[\beta_{\tau,t+L}\widehat{G}_t(y_t,\mathbf{z}_t,\mathbf{Z}_{t+1:t+L+2},\mathbf{D}_{t:t+L})]$ is $(E_{\mathbf{Z}_{\tau+1:t}}[\beta_{\tau,t}\, u_t])$-Lipschitz continuous for every $\mathbf{z}_\tau$, and the first assertion of the Lemma holds.

Our second claim has to do with the end-of-horizon cost function f_{T-L+1}. Let $1 \le \tau \le T-L+1$. The Finiteness Assumption implies that $E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L}\, f_{T-L+1}]$ is finite when $x_{T-L+1} = 0$. It is easily verified that the quantity in brackets in (5) has a derivative whose absolute value is at most $|v| + c_{T-L+1}$. Consequently, by the Finiteness Assumption, $E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L}\, f_{T-L+1}]$ is $(E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L}\, U_{T-L+1}])$-Lipschitz-continuous in x_{T-L+1}, and hence $E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L}\, f_{T-L+1}]$ is finite if x_{T-L+1} is finite. Setting $\tau = T-L+1$ we see that the second and fourth assertions of the Lemma hold for f_{T-L+1}. Also note that f_{T-L+1} is policy-independent.

Our third claim is that the order-up-to-zero policy, which we will call ζ, has finite expected cost. Let f_t^ζ and V_t^ζ be analogous to f_t and V_t in (3) and (4), corresponding to ζ. Specifically, f_t^ζ and V_t^ζ satisfy (4), and they satisfy (3) if we replace the minimum over y_t with x_t^+. We claim that $f_t^\zeta(x_t,\mathbf{z}_t)$ and $V_t^\zeta(y_t,\mathbf{z}_t)$ are finite for all finite values x_t and y_t, for all $\mathbf{z}_t$. To prove the claim note that under ζ, $y_t = (x_1 - D_{1,t-1})^+$. Hence $V_\tau^\zeta(y_\tau,\mathbf{z}_\tau) = c_\tau\, y_\tau + \sum_{t:\tau\le t\le T-L} E_{\mathbf{D}_{\tau:t+L},\mathbf{Z}_{\tau+1:t+L+2}}[\beta_{\tau,t+L}\widehat{G}_t((x_1 - D_{1,t-1})^+,\mathbf{z}_t,\mathbf{Z}_{t+1:t+L+2},\mathbf{D}_{t:t+L})] + E_{\mathbf{D}_{\tau:T-L},\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L} f_{T-L+1}((x_1 - D_{1,T-L})^+,\mathbf{Z}_{T-L+1})]$. This is finite by our first two claims. Substituting $y_t \leftarrow x_t^+$ in (3) we see that $f_t^\zeta(x_t,\mathbf{z}_t)$ is also finite for finite y_t.

We now prove the lemma by induction. For $1 \le t \le T$ we define $U_t = 2\, c_t + u_t + E[\beta_t\, U_{t+1}]$. The Finiteness Assumption implies that U_t is finite for all t. Our inductive hypothesis consists of the following affirmations.

(a) V_t is $(U_t - c_t)$-Lipschitz-continuous in y_t, the second and fourth assertions of the Lemma hold for V_t, and $V_t \le V_t^\zeta$, for $1 \le t \le T-L$ and all $\mathbf{z}_t$.

(b) f_t is U_t-Lipschitz-continuous in x_t, the third and fourth assertions of the Lemma hold for f_t, and $f_t \le f_t^\zeta$, for $1 \le t \le T-L+1$ and all $\mathbf{z}_t$.

The second claim initializes the induction by establishing Affirmation 2 for $t = T - L + 1$. Assume that we have proven Affirmation 2 for $t+1$. By (4), $V_t \leq V_t^{\zeta}$, so $V_t \in \Re \cup \{-\infty\}$ when y_t is finite. Because $U_t = 2\,c_t + u_t + E[\beta_t\, U_{t+1}]$ is finite, V_t is $(U_t - c_t)$-Lipscitz-continuous. We see that Affirmation 1 holds for t.

We now assume that Affirmation 1 holds for t and prove Affirmation 2. By (3), $f_t \leq f_t^{\zeta}$, so $V_t \in \Re \cup \{-\infty\}$ when x_t is finite. In (3), $\min_{y_t > x_t}\{K_t \delta(y_t - x_t) + V_t(y_t, \mathbf{z}_t)\} = K_t + \min_{y_t > x_t} V_t(y_t, \mathbf{z}_t)$ and $\{K_t\delta(0) + V_t(x_t, \mathbf{z}_t)\} = V_t(x_t, \mathbf{z}_t)$ are both $(U_t - c_t)$-Lipscitz-continuous in x_t, by Affirmation 2. Therefore their minimum is $(U_t - c_t)$-Lipscitz-continuous in x_t, and f_t is U_t-Lipscitz-continuous. At this point Affirmation 2 follows readily. $\diamond$

Proof of Lemma 2: Since $-f(x)$ is $[-a, -b]$-asymptotic, we can change the signs of both W and f without altering the claim. Therefore, we assume that $W \geq 0$ and $E(W) > 0$. Let $\zeta \in \{-1, 1\}$ be a constant. For every $\epsilon > 0$ and $A \geq 0$, there is a v such that for all x, $x \geq v$, we have $A + (b - \epsilon)x \leq f(x) \leq -A + (b + \epsilon)x$. By case analysis ($\zeta = 1, -1$), we see that $A \leq \zeta\,[f(x) - (b - \zeta\,\epsilon)x]$. Then

$$
\begin{aligned}
& E\{W\,\zeta\,[f(x-D) - (b - \zeta\,\epsilon)x]\} - E\{W\,\zeta\,[f(x-D) - (b-\zeta\,\epsilon)x]\,1(x - D \leq v)\} \\
= \;& E\{W\,\zeta\,[f(x-D) - (b-\zeta\,\epsilon)x]\,1(x-D > v)\} \\
= \;& E\{W\,\zeta\,[f(x-D) - (b-\zeta\,\epsilon)(x-D)]\,1(x-D > v)\} - E\{W\,\zeta\,D\,(b - \zeta\,\epsilon)\,1(x-D>v)\} \\
\geq \;& E[W A\,1(x-D>v)] - E[W\,|D|\,(|b| + \epsilon)] \\
\geq \;& A\,E[W] - A\,E[W\,1(x-D \leq v)] - (|b| + \epsilon)E[W\,|D|].
\end{aligned}
$$

Since both W and $W|D|$ have finite means, the third term is a finite constant and the second term converges to 0 as $x \to \infty$. Consider the left-hand side of the inequality. Since $f(x)$ is linearly bounded, there exist constants A' and A'' such that

$$
\begin{aligned}
& E\{|W\,\zeta\,[f(x-D) - (b-\zeta\,\epsilon)x]|\,1(x-D \leq v)\} \\
\leq \;& E\{W\,[(A' + A''|x-D|) + (|b| + \epsilon)\,x]\,1(x-D\leq v)\} \\
\leq \;& E\{W\,[(A' + A''|D|) + (A'' + |b| + \epsilon)\,x]\,1(x-D\leq v)\} \\
\leq \;& E\{W\,[(A' + A''|D|) + (A'' + |b| + \epsilon)\,(D+v)]\,1(x-D\leq v)\}.
\end{aligned}
$$

Because the expected values of W and $W\,|D|$ are both finite, the dominated convergence theorem applies, and this expression converges to 0 as $x \to \infty$.

We have proven that $\liminf_{x\to\infty} E\{W\,\zeta\,[f(x-D)-(b-\zeta\,\epsilon)x]\} \geq A\,E[W]-(|b|+\epsilon)E[W\,|D|]$. Since $E(W)>0$ and this is true for all $A\geq 0$, $E[W\,\zeta\,f(x-D)]-E[W\,\zeta\,(b-\zeta\,\epsilon)x]$ diverges as $x\to\infty$. Similarly, $E[W\,\zeta\,f(x-D)]-E[W\,\zeta\,(a-\zeta\,\epsilon)x]$ diverges as $x\to-\infty$. Considering the cases $\zeta=1$ and $\zeta=-1$, and recalling that this holds for all $\epsilon>0$, we see that $E[W\,f(x-D)]$ is $[E(W)a,\,E(W)b]$-asymptotic. $\diamond$

Proof of Lemma 3: For a given information set $\mathbf{z}_t$, suppose that whenever $x<y$ we have $V_t(x,\mathbf{z}_t)\leq V_t(y,\mathbf{z}_t)+K_t$. Then it is optimal to never order, and $s_t(\mathbf{z}_t)=-\infty$, i.e., an $(s,S)=(-\infty,S_t(\mathbf{z}_t))$ policy is optimal. Furthermore, $V_t(\cdot,\mathbf{z}_t)$ is not $[0,\cdot]$-divergent.

On the other hand, if $V_t(\cdot,\mathbf{z}_t)$ is K_t-convex and $[\cdot,0]$-divergent, and if $V_t(x,\mathbf{z}_t)>V_t(y,\mathbf{z}_t)+K_t$ for some $x<y$, then $S_t(\mathbf{z}_t)<\infty$ and the proof becomes classical. There is no local maximum $\underline{S}_t$ of V_t such that $\underline{S}_t<y$ and $V_t(\underline{S}_t,\mathbf{z}_t)>V_t(y,\mathbf{z}_t)+K_t$, because the existence of $\underline{S}_t$ would violate the K_t-convexity of V_t. Therefore $\{x: x<y \text{ and } V_t(x,\mathbf{z}_t)>V_t(y,\mathbf{z}_t)+K_t\}$ is a non-empty connected set containing $-\infty$, and V_t is non-increasing on this set. Consequently $S_t(\mathbf{z}_t)$ and $s_t(\mathbf{z}_t)$ exist, $s_t(\mathbf{z}_t)>-\infty$, and if the starting inventory level is less than $S_t(\mathbf{z}_t)$, then the $(s_t(\mathbf{z}_t),S_t(\mathbf{z}_t))$ policy is optimal. Also, we have proven that $S_t(\mathbf{z}_t)>x$ whenever $x<y$ and $V_t(x,\mathbf{z}_t)>V_t(y,\mathbf{z}_t)+K_t$. Thus, if the starting inventory level is greater than $S_t(\mathbf{z}_t)$, then it is optimal not to order; *i.e.*, the $(s_t(\mathbf{z}_t),S_t(\mathbf{z}_t))$ policy is optimal. $\diamond$

Proof of Lemma 4: We prove the lemma by induction on t starting from period $T-L$. It is easy to show that $c_{T-L}\,y$ is $[c_{T-L},c_{T-L}]$-asymptotic and $\widehat{G}_{T-L}$ is $[-\hat{\pi}_T,h_T]$-asymptotic. By Lemma 2, $f_{T-L+1}(x_{T-L+1},\mathbf{z}_{T-L+1})$ is $[E_{\mathbf{Z}_{T-L+2:T+1}}(-\beta_{T-L+1,T}\,c_{T-L+1}),E_{\mathbf{Z}_{T-L+2:T+1}}(-\beta_{T-L+1,T}\,v)]$-asymptotic in the backorder cases, and $[\cdot,E_{\mathbf{Z}_{T-L+2:T+1}}(-\beta_{T-L+1,T}\,v)]$-asymptotic in case LS. By (4), Property 1 and Lemma 2, V_{T-L} is $[\mathcal{A}_{T-L},\mathcal{B}_{T-L}]$-asymptotic ($[\cdot,\mathcal{B}_{T-L}]$-asymptotic in case LS), and the lemma holds for $t=T-L$.

Assume that the lemma holds for $t+1$, and recall that $\mathcal{B}_{t+1}>0$ by assumption. Clearly, the minimum in (3) is $[\cdot,\,\mathcal{B}_{t+1}]$-asymptotic, and f_{t+1} is $[\cdot,\,\mathcal{B}_{t+1}-c_{t+1}]$-asymptotic. In the backorder cases (PTD and PTO), if $\mathcal{A}_{t+1}<0$ then there is a finite y_{t+1} that minimizes the term in braces in (3). If $\mathcal{A}_{t+1}\geq 0$, then there may not exist a finite minimizer y_{t+1}. In either case the minimum in (3) is $[(\mathcal{A}_{t+1})^+,\,\mathcal{B}_{t+1}]$-asymptotic, and hence f_{t+1} is $[(\mathcal{A}_{t+1})^+-c_{t+1},\,\mathcal{B}_{t+1}-c_{t+1}]$-asymptotic.

By Property 1 and Lemma 2, and because $c_t\,y_t$ is $[c_t,c_t]$-asymptotic and $\widehat{G}_t$ is $[-\hat{\pi}_{t+L},h_{t+L}]$-asymptotic, V_t is $[\mathcal{A}_t,\mathcal{B}_t]$-asymptotic ($[\cdot,\mathcal{B}_t]$-asymptotic in case LS), and the lemma holds for period

t. $\diamond$

Proof of Lemma 5: We first consider the backorder cases. In these cases, the condition $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$ implies that $E_{\mathbf{Z}_{T-L+1:T+1}}[\beta_{T-L,T}(h_T + \hat{\pi}_T + c_{T-L+1} - v)] \geq 0$ and hence

$$\begin{aligned}
& V_{T-L}(y_{T-L}, \mathbf{z}_{T-L}) \\
&= c_{T-L}y_{T-L} + E_{\mathbf{D}_{T-L:T},\mathbf{Z}_{T-L+1:T+1}}[\beta_{T-L,T}\widehat{G}_{T-L}(y_{T-L}, \mathbf{z}_{T-L}, \mathbf{Z}_{T-L+1:T+1}, \mathbf{D}_{T-L:T}) \\
&\qquad\qquad + \beta_{T-L}f_{T-L+1}(y_{T-L} - D_{T-L}, \mathbf{Z}_{T-L+1})] \\
&= c_{T-L}y_{T-L} + E_{\mathbf{Z}_{T-L+1:T+1}}\{E_{\mathbf{D}_{T-L:T}}[\beta_{T-L,T}\widehat{G}_{T-L}(y_{T-L}, \mathbf{z}_{T-L}, \mathbf{Z}_{T-L+1:T+1}, \mathbf{D}_{T-L:T}) \\
&\qquad\qquad + \beta_{T-L}f_{T-L+1}(y_{T-L} - D_{T-L}, \mathbf{Z}_{T-L+1})|\mathbf{Z}_{T-L+1:T+1}]\} \\
&= c_{T-L}y_{T-L} + E_{\mathbf{Z}_{T-L+1:T+1}}\{\beta_{T-L,T}E_{\mathbf{D}_{T-L:T}}[(h_T - v)(y_{T-L} - D_{T-L,T})^+ \\
&\qquad\qquad + (\hat{\pi}_T + c_{T-L+1})(y_{T-L} - D_{T-L,T})^- - p_T D_T|\mathbf{Z}_{T-L+1:T+1}]\},
\end{aligned}$$

which is convex. Furthermore, the condition $h_{t+L} + \hat{\pi}_{t+L} \geq 0$ affirms that $\widehat{G}_t$ is convex. The proof is completed by induction on t, in the classical manner, based on (3), (4), and the condition $K_t(\mathbf{z}_t) \geq E_{\mathbf{Z}_{t+1}}[\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1})\, K_{t+1}(\mathbf{Z}_{t+1})]$.

Under lost sales, $L = 0$ and

$$\begin{aligned}
V_T(y_T, \mathbf{z}_T) &= c_T y_T + E_{D_T,\mathbf{Z}_{T+1}}[\beta_T\widehat{G}_T(y_T, \mathbf{z}_T, \mathbf{Z}_{T+1}, D_T) + \beta_T f_{T+1}((y_T - D_T)^+, \mathbf{Z}_{T+1})] \\
&= c_T y_T + E_{\mathbf{Z}_{T+1}}\{\beta_T E_{D_T}[\widehat{G}_T(y_T, \mathbf{z}_T, \mathbf{Z}_{T+1}, D_T) + f_{T+1}((y_T - D_T)^+, \mathbf{Z}_{T+1})]|\mathbf{Z}_{T+1}\} \\
&= c_T y_T + E_{\mathbf{Z}_{T+1}}\{\beta_T E_{D_T}[(h_T - v)(y_T - D_T)^+ + \hat{\pi}_T(y_T - D_T)^- - p_T D_T]|\mathbf{Z}_{T+1}\}.
\end{aligned}$$

Because $\mathcal{B}_T - \mathcal{A}_T = E_{\mathbf{Z}_{T+1}}[\beta_T(h_T + \hat{\pi}_T - v)] \geq 0$, $V_T(\cdot, \mathbf{z}_T)$ is convex.

The inductive step of the proof for the lost sales case proceeds as follows. Assume that $V_{t+1}(y_{t+1}, \mathbf{z}_{t+1})$ is K_{t+1}-convex. In (3), the classical logic implies that

$$F_{t+1}(x_{t+1}, \mathbf{z}_{t+1}) = f_{t+1}(x_{t+1}, \mathbf{z}_{t+1}) + c_{t+1}x_{t+1} = \min_{y_{t+1} \geq x_{t+1}} \{K_{t+1}\delta(y_{t+1} - x_{t+1}) + V_{t+1}(y_{t+1}, \mathbf{z}_{t+1})\}$$

is a K_{t+1}-convex function of x_{t+1}. Using the properties of $V_{t+1}(y_{t+1}, \mathbf{z}_{t+1})$ and $F_{t+1}(x_{t+1}, \mathbf{z}_{t+1})$ discussed in the second paragraph of the proof of Lemma 3, and considering the cases $s_{t+1} > 0$, $s_{t+1} = 0 < S_{t+1}$ and $S_{t+1} = 0$, we can show that $F_{t+1}(x^+_{t+1}, \mathbf{z}_{t+1})$ is a K_{t+1}-convex function of x_{t+1}, where x_{t+1} can be either positive or negative. In case LS, $\gamma(x_{t+1}) = (x_{t+1})^+$, so we can write $f_{t+1}(\gamma(x_{t+1}), \mathbf{z}_{t+1}) = -c_{t+1}x^+_{t+1} + F_{t+1}(x^+_{t+1}, \mathbf{z}_{t+1})$.

Now consider equation (4) for period t. The term in brackets is

$$\beta_t\widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1}, D_t) + \beta_t f_{t+1}((y_t - D_t)^+, \mathbf{Z}_{t+1})$$

$$= \quad \beta_t[(h_t - c_{t+1})(y_t - D_t)^+ + \hat{\pi}_t(y_t - D_t)^- + F_{t+1}((y_t - D_t)^+, \mathbf{Z}_{t+1}) - p_t D_t].$$

The condition that $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ guarantees that $(h_t - c_{t+1})(y_t - D_t)^+ + \hat{\pi}_t(y_t - D_t)^-$ is convex. Consequently, the term in brackets is K_{t+1}-convex. The condition $K_t(\mathbf{z}_t) \geq E_{\mathbf{Z}_{t+1}}[\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1})\, K_{t+1}(\mathbf{Z}_{t+1})]$ implies that the expectation of this term over $\mathbf{Z}_{t+1}$, conditioned on $\mathbf{z}_t$, is K_t-convex and the inductive step of the proof for the lost sales case is completed. $\diamond$