

Properties of the Periodic Review (R, T) Inventory Control Policy for Stationary, Stochastic Demand

Uday S. Rao

CBA, University of Cincinnati, Cincinnati, OH 45221.

February 2002, Revised September 2002, Accepted December 2002

Abstract

This paper compares the commonly-used periodic review, replenishment interval, order-up-to (R, T) policy to the continuous review, reorder point, order quantity, (Q, r) model. We show that long-run average cost function for the single product (R, T) policy has a structure similar to that of the (Q, r) model. Consequently, many of the useful properties of the latter model are applicable. In particular, the optimal cost is insensitive to the choice of the replenishment interval, T , provided the optimal order-up-to level, R , corresponding to T is used. For instance, a suboptimal T obtained from a deterministic analysis increases costs by no more than 6.125%. For continuous demand, we analytically prove that use of an (R, T) policy instead of the optimal policy increases costs by at most 41.42% in the worst-case. Computational experiments on Poisson demand demonstrate that the average-case relative error of using an (R, T) policy is under 7.5%. This relative error is lower when the demand rate and lead time are high and the fixed order costs are either very low or very high. When coordination of order placement epochs is desirable, the (R, T) policy may sometimes be preferred to the (Q, r) policy. In this context, we illustrate application of our single product results to more complex systems. In particular, we show that a simple power-of-two, (R, T) -based heuristic for the stochastic multi-product joint replenishment problem has a worst-case performance guarantee of 1.5. A similar result is explored for a special case of a two-echelon serial inventory system.

Keywords: *Inventory/Production, Joint Replenishment, Worst-Case Analysis.*

1 Introduction

In this paper, we address the classical single-item, single-stage inventory model with stationary stochastic demand (Hadley & Whitin [HW63], Silver & Peterson [SPP98], Zipkin [Z00]). In these systems, we track the inventory position, IP (= amount on-hand + on-order - backlog), over time. IP increases when we increase on-order by placing new replenishment orders and decreases when customer demands occur. Typically, inventory is controlled using one of the following ordering policies: 1. (S, s) policy: In this case, time is discretized into periods, the IP is reviewed at the beginning of each time period and, if $IP \leq s$, an order is placed for $S - IP$ units, which effectively raises IP up to S . Although this policy is optimal for the discrete-time (periodic review) problem, it has a relatively complex cost function, which complicates analysis. 2. (Q, r) policy: Here, IP is reviewed continuously and we order Q units whenever IP falls down to r . Thus, the order quantity is fixed but the time between successive orders may be random; this policy is optimal for the continuous time case. 3. (R, T) policy: IP is reviewed every T periods and, if required, an order is placed to raise the inventory position up to R . In this commonly-used policy, the minimum time between replenishment orders is fixed at T , while the order quantity is permitted to fluctuate. In all cases, there is an order-up-to or base-stock level (equal to S , $r + Q$, and R , respectively) which determines order quantities; the policies differ in the mechanism used to determine when to place orders.

The choice of which control policy to use is based on the perceived advantages of each. These advantages span the spectrum of quantitative benefits (such as lower long-run costs) and qualitative factors (such as current technological capabilities and competitive factors). For instance, managers sometimes use periodic review instead of continuous review for a whole host of reasons (e.g., ease of execution). Although continuous review is easier today than it was several years ago (due to sophisticated point-of-sale scanner systems at retail stores and automatic storage and retrieval systems at warehouses), continuous review can require expensive technological investments. Further, despite technological innovations, retail inventory status records are still sometimes not accurate enough to trigger continuous replenishments. Even when the inventory status is accurate, a periodic review may be preferred due to reduced fixed costs resulting from combining order placements for different products or due to qualitative benefits of following a regular repeating schedule of inventory replenishment. Sometimes, suppliers prefer customers to follow periodic review and replenishment

since this reduces uncertainty in the timing of demands they see. Buyers may prefer to use periodic replenishment as a means to coordinate order placements with a supplier of several items, when the supplier offers a price discount on the total dollar value of the order placed.

While the (S, s) and (Q, r) policies have been well-studied in the inventory literature, the single-item (R, T) policy has not been as well analyzed when demand is uncertain. The latter policy is convenient in many practical situations, since it allows control of the timing of order placements, which facilitates coordination of order epochs in multi-item or multi-echelon systems. As stated earlier, such coordination is commonly used to avail of quantity discounts or savings in transportation costs or operating efficiencies due to stable periodic planning. This feature of the (R, T) policy has been well exploited in deterministic systems where both order quantities, Q , and reorder intervals, T , can be kept constant (Roundy [R86], Muckstadt & Roundy [MR93]). However, the stochastic case of the (R, T) policy is yet to be similarly analyzed. In contrast, convexity and sensitivity properties of the single-item (Q, r) policy have been well studied by Zipkin [Z86] and Zheng [Z92], with extensions by Chen and Zheng [CZ94] and Gallego [G98]. Understanding of the single-item case is a crucial step towards its application to more complex systems.

In this paper, we analyze the single-item (R, T) policy and its resulting costs for different parameter settings. We assume that time is continuous, that lead time for order delivery is a known constant, and unsatisfied demand can be backlogged (backordered). There is a fixed cost for order placement; backorder and holding costs are proportional to the long-run, time-weighted average backorders and on-hand stock, respectively. To balance the trade-off between generality and analytic tractability, we model cumulative demand over an interval of time t as a stochastically increasing linear (SIL, Shaked and Shanthikumar [SS94]) process. In particular, we consider a Compound Poisson and a Brownian Motion demand process (these two distributions are described later; they capture the spectrum of stationary demand processes with independent increments). We treat both R and T as decision variables to be optimized.

The main contributions of this paper are: (1) Using novel proof techniques, we show that the long-run average cost function for the single-item (R, T) policy is jointly convex in R and T . This facilitates computation of the optimal order-up-to level R^* and the optimal reorder interval T^* , for which we present optimality conditions (newsvendor equation for R^* along with a graphical representation for T^*). Thus we extend some of the (Q, r) results from

Zipkin [Z86] and Zheng [Z92] to the (R, T) policy. (2) We demonstrate that, for continuous demand, the sensitivity results proven by Roundy [R86] and by Zheng [Z92], respectively for deterministic systems and for the (Q, r) model, extend naturally to the (R, T) model. That is, the cost of using a sub-optimal $T = \alpha T^*$, increases costs by at most $\frac{1}{2}(\alpha + \frac{1}{\alpha})$, as long as the optimal order-up-to level corresponding to the chosen T is used. Further, use of an Economic Order Interval from a deterministic analysis can provide a good approximation for the optimal T , with cost provably within 6.125% of the optimal (R, T) solution. (3) We compare cost of the periodic review (R, T) policy to a lower bound from the optimal continuous review (Q, r) policy. We develop a worst-case bound of 41.42% on the relative increase in cost of using an (R, T) policy instead of (Q, r) . While this bound is admittedly weak, it is the first worst-case performance guarantee on the commonly-used (R, T) policy. We also evaluate the empirical insensitivity and near-optimality of the (R, T) policy when demand is Poisson using numerical experiments. For the test instances from Zheng [Z92], the (R, T) policy is on average within 7.5% of optimal; its performance deteriorates when demand rates and lead times are small and when fixed order costs lie midway between a small holding cost rate and a large backorder cost rate. (4) We explore extensions relating to customer service level and illustrate application of the single-item results to the multi-product Joint Replenishment Problem (JRP, Federgruen *et al*, [FGT84]). In particular, we show that a power-of-two (R, T) based heuristic has a worst-case performance guarantee of 1.5; to the best of our knowledge, this is the first such worst-case guarantee for the stochastic JRP. Typically, such heuristics with worst-case performance guarantees exhibit significantly better average-case performance (than other heuristics or that predicted by the guarantee). Thus, we hope that our convexity and sensitivity results will provide an impetus to analyze use of the (R, T) policy in more complex multi-product, multi-stage inventory systems.

We conclude this section with a brief literature review that serves to position our work in the context of existing literature. In Section 2, we present details of our single-item model; in Section 3 we develop convexity results and optimality conditions. In Section 4, we summarize analytical sensitivity results (for Normal demand); we consider the Compound Poisson case in Section 5. We then present computational results (Section 6). In Section 7, we consider model extensions and apply our single-item results to the multi-item joint replenishment problem and a special two-echelon serial system.

1.1 Literature

The (S, s) and (Q, r) policy have been extensively studied. For instance, Zheng and Federgruen [ZF91] has considered the (S, s) policy and developed an efficient approach to compute the optimal S and s values. Our analysis of the (R, T) policy is akin to Zipkin's work on convexity ([Z86]) and Zheng's work on sensitivity ([Z92]) for the (Q, r) policy. Zipkin has shown that the long-run average cost function, $\mathcal{C}(Q, r)$, is jointly convex in Q and r ; Zhang [Z98] has provided an alternative proof of this important result. Zheng has demonstrated analytically that $\mathcal{C}(Q, r)$ is insensitive to the choice of Q whenever $r = r(Q)$, the optimal reorder point corresponding to Q . Further, Zheng provides an intuitive graphical interpretation of the optimality conditions for Q and r . More recently, Gallego [G98] has developed distribution-free bounds on Q and r and the resulting costs. In this paper, we extend some of Zipkin, Zheng, and Gallego's (Q, r) results to the (R, T) policy. We do so using proof techniques that are quite different from previous analyses of related inventory systems. For instance, we prove joint convexity of $C(R, T)$ by converting the two-dimensional cost function to a one-dimensional function and applying scalar stochastic convexity results (Shaked and Shanthikumar [SS94]). Such an approach becomes necessary because the relevant demand distribution depends on the value of decision variable T (as opposed to the (Q, r) case where the relevant distribution corresponds to demand during the fixed lead time).

Analysis of (R, T) policies has so far been restricted to deterministic demand (Roundy [R86], Muckstadt & Roundy [MR93]). However, there have been numerous diverse applications of (R, T) and the related base-stock policy (in which T is pre-specified). For instance, Hopp & Kuo [HK98] analyze and implement an (R, T) policy for cost-effective spare-parts replacement in aircraft maintenance. Encouraging computational results on (R, T) -based heuristics for joint replenishments are reported in Atkins & Iyogun [AI88]. An industrial application of base-stock policies is discussed in Rao, Scheller-Wolf and Tayur [RST00]. Our research is also similar in intent to (i) Graves ([G96]), where an inventory policy with fixed replenishment intervals has been examined computationally (for complex distribution systems consisting of a single central warehouse and multiple retailers) and (ii) Eynan and Kropp [EK98], where periodic review policies are analyzed using a simplified (approximate) cost function. Several related recent papers illustrate the potential for analysis of periodic review inventory systems, see, for instance, Cachon [C01], Moses and Seshadri [MS00] (warehouse-retailer supply chains); Graves and Willems [GW00] (setting base-stock levels); and Teunter

and Vlachos [TV01] (periodic review with flexible emergency replenishments).

2 The (R,T) Inventory Model

In this section, we describe our demand model and present the cost function.

2.1 The Demand Model

As is common practice in modeling costs for the stationary (R, T) policy, we require that demand have stationary and independent increments. This holds true for the commonly-used compound Poisson demand process and the Normal demand model (Brownian motion with drift). Let $X(t)$ be a random variable denoting cumulative demand over an interval of length t . We assume that this demand process, $X(t)$, is stochastically non-decreasing in t , with mean λt , density or mass function $\psi(x, t)$, and cumulative distribution function or CDF $\Psi(x, t)$. Further, some of our results will require that demand be continuous¹ to allow the parameters R and T to take on any non-negative, real values. We consider two specific demand models:

1. In the Brownian motion model $X(t)$ has continuous sample paths and is independent normal with mean λt , variance $\sigma^2 t$, and density

$$\psi(x, t) = \phi\left(\frac{x - \lambda t}{\sigma\sqrt{t}}\right) \equiv \frac{1}{\sigma\sqrt{2\pi t}} \exp\left[-\frac{(x - \lambda t)^2}{2\sigma^2 t}\right]. \quad (1)$$

To address the fact that this demand process is not sample-pathwise monotonic in t , we assume that there is a smallest time unit, say $t = t_{min} \equiv 1$, and that all feasible reorder intervals and lead time data must be at least t_{min} . Further, demand rate λ is sufficiently larger than σ (e.g., $\lambda > 3.5\sigma$), so that $\lambda t \gg \sigma\sqrt{t}$ for all $t \geq t_{min}$ and, consequently, the probability of negative demand is negligibly small for $T \geq t_{min}$.

2. The compound Poisson demand model is a jump process with

$$\psi(x, t) \equiv P(X(t) \approx x) = P\left(\sum_{i=0}^{N(t)} Y_i \approx x\right) = \sum_{n \geq 0} P\left(\sum_{i=0}^n Y_i \approx x\right) \cdot P(N(t) = n), \quad (2)$$

where (i) the Y_i denote independent, identically distributed, non-negative, random variables corresponding to the quantity demanded at a demand epoch; (ii) $N(t)$ is

¹We note that the only monotonic stochastic process with stationary, independent increments *and* continuous sample paths is the deterministic process, $X(t) = \lambda t$ (see Feller [F71], Chapter IX). Consequently, the continuous model must be viewed as a convenient tool for approximating the real system.

a Poisson point process, independent of any Y_i , with parameter ν and $P(N(t) = n) = e^{-\nu t}(\nu t)^n/n!$, that characterizes demand epochs, and (iii) $P(\sum_{i=0}^n Y_i \approx x)$ represents the n -fold convolution of the probability density (or mass) function of Y_i .

Our analyses in Sections 3 and 4 will focus on the Brownian motion model, the compound Poisson case will be discussed briefly in Section 5. In general, we restrict attention to the case when $X(t)$ is SIL (Stochastically Increasing and Linear). For a summary of stochastic convexity and its implications refer to Section 5.3 of Chang, Shanthikumar and Yao [CSY94] or Chapter 6 in Shaked and Shanthikumar [SS94]. For completeness, we include the following definitions:

Definition 1 *If $Z(\theta)$ is a random variable whose CDF is parameterized by θ , then $Z(\theta)$ is SICX (Stochastically Increasing and ConveX) in θ if, for any ICX function $f(z)$, $E[f(Z(\theta))]$ is ICX in θ . The terms SI and SCX are defined analogously. $Z(\theta)$ is SIL (Stochastically Increasing and Linear) if it is both SICX and SICV (where CV = ConcaVe).*

Using the above definition, it is easy to verify the following properties², which we will use to prove convexity of the long-run average cost function.

Proposition 2 ([SS88]) *The compound Poisson demand process $X(t)$ is SIL in t .*

Proposition 3 *When $\lambda t_{min} > 3.5\sigma\sqrt{t_{min}}$, the Brownian motion process $X(t)$ is SIL, for $t \geq t_{min}$.*

2.2 The Cost Function $C(R, T)$

Let L denote the fixed lead time for order delivery from a perfectly reliable, external supplier. The cost parameters are: K , the fixed order placement or setup cost; h , the inventory holding cost rate; and p , the backorder cost rate. Under the general conditions on demand described in Serfozo & Stidham [SS78], the long-run average cost function for the (R, T) policy is:

$$C(R, T) = \frac{K}{T}P(X(T) > 0) + H(R, T), \quad \text{with } H(R, T) = \frac{1}{T} \int_L^{L+T} G(R, t)dt, \quad (3)$$

where, following Zheng [Z92], we have pooled the inventory holding and backordering costs into the function G . That is, for $x^+ = \max(0, x)$,

$$G(R, t) = hE[(R - X(t))^+] + pE[(X(t) - R)^+] \quad (4)$$

²Missing proofs are either included in the appendices or are available from the author.

$$= h(R - \lambda t) + (h + p) \int_{x=R}^{\infty} (x - R)\psi(x, t)dx \quad (5)$$

$$= p(\lambda t - R) + (h + p) \int_{x \leq R} \Psi(x, t)dx, \quad (6)$$

where we have used the fact that $x^+ = x + (-x)^+$ and $E[(R - X)^+] = \int_{x \leq R} P(X \leq x)dx$.

The first term in (3) represents the long-run average setup costs. When demand is continuous (e.g., a Brownian motion process with positive drift), $P(X(T) > 0) = 1$ for any feasible T . When demand is compound Poisson, $P(X(T) > 0) = 1 - e^{-\eta T}$, where $\eta = \nu[1 - P(Y_i = 0)]$ is a positive constant. In either case, the long-run setup cost portion, $\frac{K}{T}P(X(T) > 0)$, is decreasing and convex³ in T . For simplicity, we set $P(X(T) > 0) \equiv 1$, with the understanding that all our analytical results may be easily modified to apply to the case when $P(X(T) > 0) = 1 - e^{-\eta T}$. In the next section, we show that the long-run average inventory cost term $H(R, T)$ is convex.

3 Convexity and the Optimality Conditions

In this section, we prove that $C(R, T)$ is jointly convex in R and T , when demand $X(t)$ is SIL. We also develop equations that characterize the optimal values of R and T ; in particular, we provide a newsvendor equation for $R(T)$, the optimal R given T , and present the modified demand process to be used in this newsboy calculation. These equations allow us to prove sensitivity properties and to efficiently compute the optimal parameter values, R^* and T^* . The solution methodology we follow may be summarized as follows: We use a transformation to decompose the cost function into more manageable pieces. Specifically, for a given reorder interval, we define a modified demand process that represents the demand over the reorder interval plus a random portion of the lead time. Then, part of the cost function resembles the traditional newsvendor objective function and standard arguments can then be used to establish convexity in R . Further, it follows that, given a reorder interval T , we can express the optimal base-stock level $R(T)$ using a critical fractile solution. Given $R(T)$, the optimal reorder interval T^* can be obtained by solving an EOQ-type of problem (with a fixed cost portion decreasing in T and an inventory cost portion that increases with T).

³For the compound Poisson case, the second derivative of $\frac{K}{T}P(X(T) > 0)$ is $\frac{K}{T^3}[2 - 2e^{-\eta T} - 2\eta T e^{-\eta T} - \eta^2 T^2 e^{-\eta T}] \equiv \frac{K}{T^3}[f(T)]$. Since $f(0) = 0$ and $f'(T) = \eta^3 T^2 e^{-\eta T} \geq 0$, we infer that $f(T) \geq 0$ for all $T \geq 0$, which implies convexity of $\frac{K}{T}P(X(T) > 0)$.

Combining equations (3) and (5), we see that

$$H(R, T) \equiv h(R - E[Y(T)]) + (h + p)E[(Y(T) - R)^+], \quad (7)$$

where $Y(T)$ has density $\gamma(x, T) = \frac{1}{T} \int_L^{L+T} \psi(x, t) dt$. Thus, $Y(T) \stackrel{d}{=} X(L + \mathcal{U}(0, 1)T)$, where the equality is in distribution and $\mathcal{U}(a, b)$ denotes a uniform random variable between a and b . Given T , $H(R, T)$ has the form of a standard newsvendor average cost function with $Y(T)$ denoting the relevant demand random variable. Note that traditional formulas for the order-up-to level in periodic review inventory systems approximate inventory costs by effectively computing costs at the end of each replenishment cycle. Hence, they use demand over the “protection interval” $L + T$, which is $X(L + T)$ in our notation. We compute inventory costs continuously within each cycle; a typical cycle starts L time units after an order is placed (i.e., at $t = L$) and ends T time units later (i.e., $t = L + T$). Thus, the observed demand during the cycle varies between $X(L)$ and $X(L + T)$, and $X(L + \mathcal{U}(0, 1)T)$ is the appropriate equivalent demand that should be used to compute the order-up-to level corresponding to the newsvendor critical ratio. Since we consider both R and T as decision variables to be optimized, we need to understand how $C(R, T)$ changes with R and T ; joint convexity of $C(R, T)$ would simplify computation of the optimal R^* and T^* .

Note that $C(R, T)$ is convex if $H(R, T)$ is convex. Further, $E[Y(T)] = E[\lambda(L + \mathcal{U}(0, 1)T)] = \lambda L + \frac{1}{2}\lambda T$, which is linear in T . Consequently, by (7), $C(R, T)$ is convex if $B(R, T) = E[(Y(T) - R)^+]$ is convex. To prove convexity of $B(R, T)$ we use the following easily proven result (Rockafeller, [R70]):

Fact 4 *Let V denote the vector (R, T) . Then $B(V)$ is convex in V if and only if, for each $V_i = (R_i, T_i)$, $i = x, y$, the function $b(\alpha) = B(\alpha V_x + (1 - \alpha)V_y)$ is convex in α for $0 \leq \alpha \leq 1$.*

This result allows us to convert joint convexity in the two-dimensional vector (R, T) to convexity in a single variable, α . To show convexity of $b(\alpha)$, we use single (scaler) parameter, stochastic convexity results developed by Shaked and Shanthikumar [SS88]. The following lemma, proven in Appendix A, forms the basis of our proof of convexity of $C(R, T)$.

Lemma 5 *Let $\alpha \in [0, 1]$. (i) For any $R_x, R_y \geq 0$, $\tilde{R}(\alpha) \equiv \alpha R_x + (1 - \alpha)R_y$ is linear in α , hence, $\pm \tilde{R}(\alpha)$ is SCX in α . (ii) $Y(T) \equiv X(L + \mathcal{U}(0, 1)T)$ is SIL in T . (iii) For any $T_x, T_y \geq 0$, $\tilde{Y}(\alpha) \equiv Y(\alpha T_x + (1 - \alpha)T_y)$ is SCX in α . (iv) $\tilde{Y}(\alpha) - \tilde{R}(\alpha)$ is SCX in α .*

Theorem 6 *$C(R, T)$ is jointly convex in R and T .*

Proof. By Lemma 5 (iv), $\tilde{Y}(\alpha) - \tilde{R}(\alpha)$ is SCX in α . Consequently, by Definition 1 and the convexity of x^+ , $E[(\tilde{Y}(\alpha) - \tilde{R}(\alpha))^+] = b(\alpha)$ is convex in α . Hence, by Fact 4, $B(R, T)$ is jointly convex, which yields the required result. \blacksquare

By the above theorem, the optimal R^* and T^* may be found easily using standard approaches for convex minimization (such as golden section search, refer to Bazaraa *et al* [BSS93]). In the remainder of this section we develop optimality conditions for R and T .

3.1 The Optimal R , Given T

Let $R(T) = \operatorname{argmin}_R C(R, T)$ denote the optimal R corresponding to reorder interval T . By equation (7), it follows that $R(T)$ may be obtained by solving a “newsvendor” problem with demand $Y(T)$ and cost $H(R, T)$. That is, for continuous demand, $R(T)$ satisfies, $\partial H(R, T)/\partial R = 0$ or

$$hP(Y(T) \leq R(T)) = pP(Y(T) \geq R(T)) = \frac{hp}{h+p}. \quad (8)$$

By equation (8), since $Y(T)$ is stochastically increasing in T , $R(T)$ is non-decreasing in T . In fact, by differentiating equation (8), we see that

$$R'(T) = \frac{P(X(L+T) \geq R(T)) - P(Y(T) \geq R(T))}{T\gamma(R(T), T)} > 0 \text{ for all } T > 0, \quad (9)$$

where $\gamma(\cdot)$ denotes the density of $Y(T)$. Thus $R(T)$ is increasing in T and $R(T) \rightarrow \infty$ as $T \rightarrow \infty$. To facilitate the sensitivity analysis in Section 4, it is useful to identify bounds on $\lim_{T \rightarrow \infty} R'(T)$. In this context, we note that, when $E[X(t)] = \lambda t$, it is possible to argue that the denominator in (9) satisfies $\lim_{T \rightarrow \infty} T\gamma(x, T) \equiv \lim_{T \rightarrow \infty} \int_L^{L+T} \psi(x, t) dt \leq \frac{1}{\lambda}$. In fact, by equation 9 in Appendix 4 of [HW63], we see that for Normal demand,

$$\gamma(x, T) \approx \frac{1}{\lambda T} \left[\bar{\Phi} \left(\frac{x - \lambda(L+T)}{\sigma\sqrt{L+T}} \right) - \bar{\Phi} \left(\frac{x - \lambda(L)}{\sigma\sqrt{L}} \right) \right], \quad (10)$$

where $\bar{\Phi}(x)$ is the complementary cumulative distribution function (CCDF) for the standard normal. In this case, $\lim_{T \rightarrow \infty} T\gamma(R(T), T) = \lim_{T \rightarrow \infty} \frac{1}{\lambda} \bar{\Phi}(w_T)$, where $w_T = \frac{R(T) - \lambda(L+T)}{\sigma\sqrt{L+T}}$. Combining this with equations (8) and (9), we get:

$$\lim_{T \rightarrow \infty} R'(T) = \lim_{T \rightarrow \infty} \frac{\bar{\Phi}(w_T) - \frac{h}{h+p}}{\frac{1}{\lambda} \bar{\Phi}(w_T)} = \lambda \left[1 - \frac{h}{(h+p) \lim_{T \rightarrow \infty} \bar{\Phi}(w_T)} \right]. \quad (11)$$

The right hand side in (11) is bounded below by 0 and above by $p/(p+h)$.

3.2 The Optimal T (Continuous Demand)

In this section we present an optimality condition for T . This condition is given a graphical interpretation which turns out to be useful in developing the analytical sensitivity results and in explaining the computational observations.

Define $C(T) = C(R(T), T)$ and $U(T) = \int_L^{L+T} G(R(T), t)dt$ so that $C(T) = \frac{K+U(T)}{T}$.

Proposition 7 *For continuous demand, the cost function $C(T)$ can be written in the form: $C(T) = \frac{K}{T} + \frac{1}{T} \int_0^T u(\tau)d\tau$, where $u(\tau) = G(R(\tau), L + \tau)$.*

Proof. Observe that $U(0) = 0$, $U(T)$ is differentiable and $U'(T)$ is continuous. Consequently, using elementary calculus, $U(T) = \int_0^T u(\tau)d\tau$, where $u(\tau) = U'(\tau) = G(R(\tau), L + \tau) + R'(\tau) \int_L^{L+\tau} \frac{\partial}{\partial R} G(R(\tau), t)dt = G(R(\tau), L + \tau)$, since the latter integral is zero. ■

Proposition 8 *$C(T)$ is convex. If T^* minimizes $C(T)$, then $C(T^*) = u(T^*)$ and hence, by Proposition 7, $K + \int_0^{T^*} u(\tau)d\tau = T^*u(T^*)$.*

Proof. Convexity of $C(T)$ follows from the joint convexity of $C(R, T)$ and a convexity preservation theorem (see, e.g., Proposition B-4, page 525, in Heyman and Sobel [HS84]). By Proposition 7, $C'(T) = -\frac{K}{T^2} - \frac{\int_0^T u(\tau)d\tau}{T^2} + \frac{u(T)}{T} = -\frac{C(T)}{T} + \frac{u(T)}{T}$. Hence, at $T^* = \operatorname{argmin} C(T)$, we have $C'(T^*) = 0$ or $C(T^*) = u(T^*)$. ■

Proposition 9 *$u(\tau)$ is a non-decreasing, convex function.*

Interpretation of Propositions 7–9: By the convexity of $C(T)$ in Proposition 8, the optimal T may be found using a simple one-dimensional golden section search in T ; the optimal cost is $C(T^*) = u(T^*)$. By Proposition 7, we interpret $u(\tau)$ as the rate at which expected inventory costs accumulate at time $L + \tau$, if the order-up-to level is $R(\tau)$. This yields an interesting graphical interpretation of the optimality condition for T in Proposition 8 which may be illustrated as follows: A typical graph of $u(T)$, for continuous demand, is shown in Figure 1, where as per Proposition 8, $u(T^*) = C(T^*)$ and the area under $u(\cdot)$ between 0 and T measures the inventory cost per cycle. The total cost per cycle is given by the rectangular area $u(T^*)T^*$ and is comprised of a setup cost portion, K , which lies above $u(T)$, and the inventory cost per cycle.

Note that the form of $C(T)$ in Proposition 7 is directly applicable to certain spare parts replenishment problems. In this case, the function $u(\tau)$ denotes the cost (or system performance loss) per unit time incurred by a component at age τ , and $\frac{1}{T}$ specifies the component replacement frequency. See ([HK98]) for details.

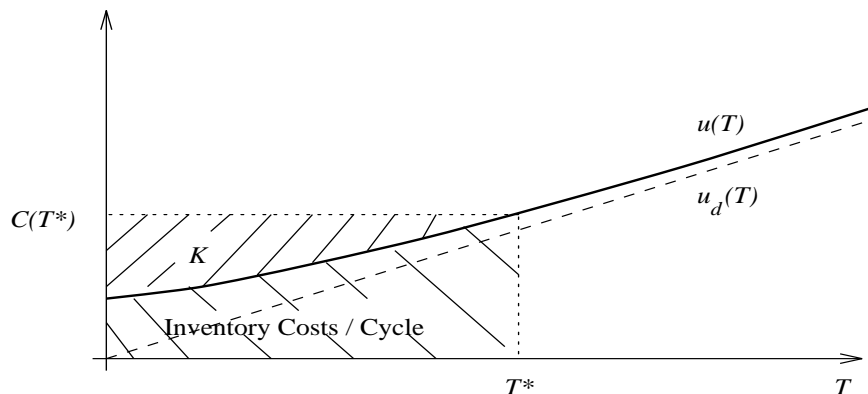


Figure 1: Graph of $u(T)$ showing T^* and $C(T^*)$

We conclude this section by defining the function $u_d(T)$ in Figure 1. Let $u_d(T)$ denote the inventory cost rate function, $u(T)$, for the deterministic demand model. From Hadley & Whitin [HW63], we know that, with $H \equiv \frac{hp}{h+p}$,

$$C_d(T) = \frac{K}{T} + \int_0^T u_d(\tau) d\tau \text{ for } u_d(T) = \frac{hp\lambda T}{h+p} = H\lambda T, \text{ with } u'_d(T) = H\lambda. \quad (12)$$

By the convexity of $u(T)$ as a function of $X(L+T)$ and T , and by Jensen's inequality,

$$u(T) \geq u_d(T) = H\lambda T. \quad (13)$$

Hence, in Figure 1, the line $u_d(T)$ is below the function $u(T)$; the fact that $u_d(T)$ is parallel to $u(T)$ as $T \rightarrow \infty$ will follow from Proposition 10 stated below.

4 Sensitivity Properties of $C(\mathbf{R}, \mathbf{T})$

In this section we focus on the continuous demand case; the discrete demand case is explored in Sections 5 and 6. All sensitivity properties of interest follow from Proposition 7 and the following result, proven in Appendix A:

Proposition 10 *If $H \equiv \frac{hp}{h+p}$ and $u(T)$ is convex, then $u'(T) \leq \lim_{T \rightarrow \infty} u'(T) = H\lambda$.*

We use the above result to extend the sensitivity properties from Roundy [R86] and Zheng [Z92] to the stochastic (R, T) inventory model. Given the structure of $C(T)$, the following result follows directly from arguments similar to those used by Zheng [Z92] to prove the corresponding result for the (Q, r) model.

Theorem 11 ((R, T) counterpart of Roundy [R86] & Zheng [Z92]) *If T^* minimizes $C(T)$ then, for any $\alpha > 0$,*

$$\frac{C(\alpha T^*)}{C(T^*)} \leq \frac{1}{2}\left(\alpha + \frac{1}{\alpha}\right)$$

Thus, $C(T)$ is insensitive to the choice of T : If $\frac{1}{2}T^ \leq T \leq 2T^*$, then $C(T) \leq 1.25C(T^*)$; rounding T^* to a power-of-two multiple corresponds to $\alpha = \sqrt{2}$ and $C(T) \leq 1.06C(T^*)$.*

Proof. $\frac{C(\alpha T^*)}{C(T^*)} = \frac{K + \int_0^{\alpha T^*} u(\tau) d\tau}{\alpha T^* C(T^*)} = \frac{1}{\alpha} \left[1 + \frac{\int_{T^*}^{\alpha T^*} u(\tau) d\tau}{T^* C(T^*)} \right] \stackrel{\text{Lemma 18}}{\leq} \frac{1}{\alpha} \left[1 + \frac{\frac{1}{2}(\alpha^2 - 1)u(T^*)}{C(T^*)} \right]$, which yields the desired result since $u(T^*) = C(T^*)$. ■

Let T_d^* denote the optimal solution to a deterministic Economic Order Interval problem. From Hadley & Whitin [HW63],

$$T_d^* = \sqrt{\frac{2K}{H\lambda}}, \quad R_d^* = \lambda L + \frac{p\lambda T_d^*}{h + p}. \quad (14)$$

Under stochastic demand, we compare the (R, T) policy having optimal parameter values with an $(R(\sqrt{2}T_d^*), \sqrt{2}T_d^*)$ policy, where $R(\sqrt{2}T_d^*)$ solves equation (8). Since $u_d(T) \leq u(T)$ and $0 \leq u'(T) \leq u'_d(T)$, as per Figure 1, $u(T)$ requires a larger T to trap an area K above it than does $u_d(T)$. Hence, $T^* \geq T_d^*$. Further, by the monotonicity of $u(\tau)$, $\frac{1}{T} \int_0^T u(\tau) d\tau$ is increasing in T , consequently, $\frac{1}{T^*} \int_0^{T^*} u(\tau) d\tau \geq \frac{1}{T_d^*} \int_0^{T_d^*} u(\tau) d\tau$. This yields the following result, the proof of which is analogous to Gallego's [G98] proof of the corresponding result for the (Q, r) policy.

Theorem 12 *If T^* minimizes $C(T)$, and T_d^* is the optimal deterministic order interval defined by equation (14) then, $T^* \geq T_d^*$ and, as in Zheng [Z92] and Gallego [G98], resp.:*

$$\frac{C(T_d^*) - C(T^*)}{C(T^*)} \leq 12.5\% \quad \text{and} \quad \frac{C(\sqrt{2}T_d^*) - C(T^*)}{C(T^*)} \leq 6.125\%$$

Thus, use of a reorder interval obtained from a deterministic analysis increases costs by at most 6.125%.

In the remainder of this section, we bound the relative increase in cost from using the (R, T) policy instead of the optimal policy.

Lemma 13 (Gallego [G98]) *For the (Q, r) model with $\mathcal{G}(y) \equiv E[h(y - X(L))^+ + p(X(L) - y)^+]$, $\mathcal{G}(y^*) \equiv \min_y \mathcal{G}(y)$, and long-run average cost $\mathcal{C}(Q, r) = \frac{1}{Q}[\lambda K + \int_r^{r+Q} \mathcal{G}(y) d\tau]$, we have*

$$\mathcal{C}(Q^*) = \mathcal{C}(Q^*, r(Q^*)) \geq \sqrt{2\lambda KH + \mathcal{G}(y^*)^2}.$$

Lemma 14 *Let $\mathcal{G}(y^*)$ be as defined in Lemma 13. Then*

$$C(T^*) = C(R(T^*), T^*) \leq \sqrt{2\lambda KH} + \mathcal{G}(y^*)$$

Proof. Consider the decomposition $C(T) = C_0(T) + u(0)$, where $C_0(T) = \frac{1}{T}[K + \int_0^T u_0(\tau)d\tau]$ with $u_0(\tau) = u(\tau) - u(0)$. Then $T^*C_0(T^*) = T^*u_0(T^*) \leq 2K = T_d^*\sqrt{2\lambda KH}$, where the inequality follows from the convexity of $u_0(\tau)$. (From Figure 1, we see that the rectangular area $T^*[C(T^*) - u(0)]$ is no more than $2K$.) Thus $C_0(T^*) \leq \frac{T_d^*}{T^*}\sqrt{2\lambda KH}$. Since $T_d^* \leq T^*$, $C(T^*) \leq \sqrt{2\lambda KH} + u(0)$. Now the result follows from $u(0) = G(R(0), L) = \mathcal{G}(y^*)$. ■

Note that the $u(0)$ in the above lemma is a constant which corresponds to the “newsvendor” cost resulting exclusively from variability in demand. This cost is incurred even for $K = 0$, in which case we place orders as often as required.

Theorem 15 *For the continuous demand model, in the worst-case,*

$$\frac{C(R(T^*), T^*) - \mathcal{C}(Q^*, r(Q^*))}{\mathcal{C}(Q^*, r(Q^*))} \leq \sqrt{2} - 1.$$

That is, use of an (R, T) policy with optimal parameters can cost at most 41.42 % more than use of a (Q, r) policy.

Proof. Observe that $C(R(T^*), T^*) \leq \sqrt{2\lambda KH} + \mathcal{G}(y^*) \leq \sqrt{2} \sqrt{2\lambda KH + \mathcal{G}(y^*)^2} \leq 1.4142 \mathcal{C}(Q^*, r(Q^*))$ ■

Note that a weak worst-case result such as Theorem 15 may be of little practical significance. However, in practice, an (R, T) policy may result in K, h values which are effectively smaller than those for the corresponding (Q, r) policy due to reduced costs of tracking inventory and due to coordination of order placements among different items (thereby sharing some fixed order costs in a multi-product setting). Any such reduction in the parameters K and h would reduce the increase in cost of using an (R, T) policy. Furthermore, even if the K and h cost parameters remain identical in implementing the (Q, r) and (R, T) policies, our computational experimentation suggests that the actual difference between optimal costs of these two policies is, on average, significantly smaller than the worst-case bound of 41.4% (though instances nearly achieving this bound do exist). Before we present details of our computational results, we consider the compound Poisson demand model.

5 Compound Poisson Demand

Besides the brownian motion model considered above, our computational testing will use discrete Poisson demand (problem instances from Zheng [Z92]). Hence, in this section, we briefly explore the Compound Poisson demand model in which the quantity demanded (and hence, the optimal R) could be discrete. First, we note that the convexity results from Section 3 remain valid since we did not use continuity of R and T in proving convexity of $C(R, T)$ and $C(T)$. Thus, the optimal values of R and T may be found using convex minimization techniques, such as the golden section method (which only requires $C(T)$ to be unimodal, and not necessarily smooth or convex). For the general compound Poisson model, one problem in implementing such a search is that evaluation of $C(T)$ for fixed T may require simulation or complex convolutions to determine inventory costs $H(R, T)$. Hence, in our computational testing, we restrict attention to the simple Poisson case, which allows for comparison with the (Q, r) results in Zheng [Z92]. When demand is discrete, $R(T)$ must be determined using the discrete version of the newsvendor: That is, $R(T)$ is the smallest value of R satisfying $P(Y(T) \leq R) \geq \frac{p}{h+p}$. Since equality may not hold in equation (8), Proposition 7 does not apply. Yet, we conjecture that sensitivity results similar to those for the normal demand model may apply for most commonly used Y_i (discrete or continuous). However, the loss of differentiability for general Y_i makes rigorous analysis a challenging area for future research, particularly for the discrete time case (where the (R, T) policy would need to be compared to the (S, s) policy). A starting point for such an analysis may be obtained from Section 3 of Zheng and Chen [ZC92] where the continuous-time (nQ, r) policy is considered when the quantity demanded is discrete.

6 Computational Results

For consistency and comparability, the set of test problems we use are identical to the 135 problems used in Zheng, [Z92]. In all problem instances, the lead time is normalized to $L = 1$. Demand is Poisson with rate λ taking on values 5, 25 and 50; K is 1, 5, 25, 100, or 1000; h is 1, 10 or 25; p is 5, 10, 25 or 100. As is common in practice (Zheng, [Z92]), we restrict attention to $p \geq h$. Note that Zheng's data set contains 45 problem instances with $p < h$. In our experimentation on these problems, results similar to the case of $p \geq h$ were obtained. Further, use of normal demand instead of Poisson demand did not significantly alter the

results. Table 1 lists the computed parameter values and costs for three typical numerical problem instances which, for each K value, differ only in the h/p ratio.

Table 1: Some Numerical Examples: Policy Parameters and Costs

$$\%Dev1 = \frac{C(R(T_d^*), T_d^*) - C(R(T^*), T^*)}{C(R(T^*), T^*)} \times 100; \quad \%Dev2 = \frac{C(R(T^*), T^*) - C(r^*, Q^*)}{C(r^*, Q^*)} \times 100$$

Poisson Demand, $\lambda = 50, L = 1, h = 10, p = 25$.									
K	T_d^*	T^*	$R(T_d^*)$	$R(T^*)$	r^*	Q^*	$C(R^*, T^*)$	$\%Dev1$.	$\%Dev2$.
1	0.075	0.139	56.0	58.0	50.0	7.0	98.56	2.40	3.25
5	0.167	0.226	58.0	60.0	48.0	12.0	120.43	2.41	4.44
25	0.374	0.446	65.0	67.0	44.0	23.0	179.39	0.98	4.85
100	0.748	0.806	77.0	79.0	38.0	40.0	300.45	0.17	3.83
1000	2.366	2.392	134.0	135.0	15.0	120.0	864.97	0.00	1.46
Poisson Demand, $\lambda = 50, L = 1, h = 20, p = 20$									
K	T_d^*	T^*	$R(T_d^*)$	$R(T^*)$	r^*	Q^*	$C(R^*, T^*)$	$\%Dev1$.	$\%Dev2$.
1	0.063	0.118	51.0	53.0	46.0	7.0	127.56	2.95	2.66
5	0.141	0.209	53.0	55.0	44.0	10.0	151.80	3.09	3.37
25	0.316	0.381	58.0	59.0	39.0	20.0	217.65	1.60	3.86
100	0.632	0.693	65.0	67.0	32.0	35.0	355.92	0.36	2.81
1000	2.000	2.020	100.0	100.0	-1.0	103.0	1019.90	0.01	0.89
Poisson Demand, $\lambda = 50, L = 1, h = 15, p = 100$									
K	T_d^*	T^*	$R(T_d^*)$	$R(T^*)$	r^*	Q^*	$C(R^*, T^*)$	$\%Dev1$.	$\%Dev2$.
1	0.055	0.100	60.0	61.0	56.0	4.0	194.76	1.92	1.74
5	0.124	0.172	62.0	63.0	53.0	10.0	223.77	2.01	4.25
25	0.277	0.339	67.0	69.0	50.0	17.0	301.72	1.01	5.37
100	0.554	0.583	77.0	78.0	47.0	32.0	462.44	0.21	5.95
1000	1.751	1.798	127.0	129.0	38.0	91.0	1217.31	0.01	3.68

In Table 1, $\%Dev1$ represents the percentage increase in cost due to “*error in optimization of T* ”, i.e., due to incorrectly choosing $T^* = T_d^*$; $\%Dev2$ is the percentage relative “*error in policy*” which results from use of an (R, T) policy instead of a (Q, r) policy. In general, the relative “error in optimization of T ” is small (at most 3.1%). The relative difference between using an (R, T) policy and a (Q, r) policy is more pronounced, but still within six percent for the tabulated problem instances. Further experimentation has revealed that the cost of using an (R, T) policy is relatively insensitive to the T value (rounding to a power-of-two multiple of a base planning period did not increase costs substantially). However, it is more important to choose the optimal R for the specified T .

For the data in Table 1, $\%Dev1$ is highest for $K = 5$ and $\%Dev2$ is highest when $K = 25$. Our computational testing suggests that both $\%Dev1$ and $\%Dev2$ become smaller for very

small or very large K . This may be explained as follows: By Figure 1 we see that for large K , $T_d^* \approx T^*$ so %Dev1 is small. Further, the $\sqrt{2\lambda KH}$ portion of total costs dominates the inventory costs due to uncertainty. Consequently $C(R(T^*), T^*)$ is closer to $C(r^*, Q^*)$ and %Dev2 is small. For the case of small K , T^* is small, the inventory is reviewed frequently and hence the error resulting from periodic review (%Dev2) is small. While the ratio T^*/T_d^* may be large, the curve $u(T)$ in Figure 1 is quite flat at low T values, as a result the cost function is not too sensitive to T and consequently %Dev1 is small. Along similar lines, we noted that, when compared to %Dev1, a larger increase in K is required to make %Dev2 “small”. In addition, we observed that:

1. Whether demand is modeled using a continuous distribution or a discrete distribution, it is easy to compute the optimal R and T using a simple search technique such as the golden section method. Requisite upper bounds on T^* may be obtained using $H\lambda T^* \leq C(R^*, T^*) \leq \mathcal{M}$, where \mathcal{M} is an easily computed demand distribution-free ([G98]) upper bound on costs.
2. As expected, T and $R(T)$ increase when K or p increases or when h decreases. However, no discernible monotonicity in parameters or costs resulted from increasing the h/p ratio. In general, $R(T^*)$ is close to $r^* + Q^*$ and T^* is close to Q^*/λ . The relative difference in parameter values, $T_d^* \& T^*$ and $R(T_d^*) \& R(T^*)$, decreases as K increases. (This is consistent with Figure 1 which indicates that $T^* \approx T_d^*$ for large K . Note that for the Poisson demand model $u(T)$ is non-decreasing as in Figure 1, however, since R takes on discrete values, $u(T)$ may not be smooth.)
3. The actual gap in long-run average costs between using the optimal T^* and the sub-optimal T_d^* is much smaller in practice than the worst-case bound of 12.5%. In all cases, an optimal order-up-to level, $R(T)$, corresponding to the reorder interval T must be used. (Note that although use of a reorder interval of $\sqrt{2}T_d^*$ results in an improved worst-case performance of 6.125 %, in practice $\sqrt{2}T_d^*$ may be further⁴ from optimal than T_d^* , particularly when relatively large values of K make $T^* \approx T_d^*$.)
4. The optimal cost of using the (R, T) policy is often reasonably close to the optimal cost of a (Q, r) policy. From our computational experiments (on different sets of test

⁴It is possible to modify the $\sqrt{2}T_d^*$ heuristic to circumvent this problem; in fact, a better demand distribution-free reorder interval may also be determined.

problems), it appears that the performance of the (R, T) policy deteriorates for small values of λ and L and $h \ll K \ll p$, where \ll denotes “less than by an order of magnitude.” For instance, with $h = 1$, $K = 10$, $p = 100$, $\lambda = 1$, $L = 1$, the (R^*, T^*) policy was almost 41 % worse than a (Q^*, r^*) policy.

An explanation for this is as follows: When the demand rate λ is small, demand is sporadic, and this favors use of a continuous review policy. When lead time L is small, the (Q, r) policy, whose optimal parameters Q^* and r^* are based on demand during lead time, is more efficient in that it can manage with very little inventory (and still keep the duration of the rare shortage small). On the other hand, the inventory carried by the (R, T) policy depends on both L and T , and this policy can incur higher inventory and shortage costs when T^* is larger than 0. We have already argued earlier that for very small or very large K the (R, T) policy will work effectively. When the holding cost is relatively low and the shortage cost is high, the (R, T) policy is more likely to fall into the trap of carrying, and paying for, more inventory than the (Q, r) policy. However, under sporadic demand, there are scenarios in which this extra inventory is not sufficient to avoid stockouts and this results in increased shortage penalty costs relative to the (Q, r) policy. Under these circumstances, the increased inventory carrying and shortage costs of the (R, T) policy are not sufficiently offset by the relatively small savings in fixed order costs.

Results for the 135 problems tested may be summarized as follows: The error in optimization of T , indicated by the percent deviation, $\%Dev1$, was within 0.25% for over 50% of the problems and between 0.25% and 0.75% for 20% of the problems. The average value of $\%Dev1$ was 0.63%, the maximum was 3.1%. As expected, the relative error in policy was much higher with $\%Dev2$, having min, mean and max values of 0.36%, 7.31% and 23.32% respectively. Overall, the computational testing supports the conclusion that, in practical multi-item, multi-echelon systems where coordination of order placement epochs is critical, an (R, T) policy may be a viable alternative for controlling inventory of certain classes of high demand items.

7 Discussion, Extensions and Future Research

7.1 Customer Service Requirements

If, instead of specifying a penalty cost rate p , a Type I cycle service level of α is prescribed, then the optimal $R(T)$ may be obtained using equation (8) with $p = \frac{\alpha h}{1-\alpha}$. That is, we solve $P(Y(T) \leq R(T)) = \alpha$ for $R(T)$. Since the cycle service requirement does not change the structure of the problem, the results from the previous sections are still applicable.

If $p = 0$ and a Type II fill rate of β is prescribed, then $R(T)$ is obtained using $E[(Y(T) - R)^+] = (1 - \beta)\lambda T$. Since $E[(Y(T) - R)^+]$ is convex in R , the optimal R for each T may be obtained using a golden section search. Further, the long-run average cost $C(R, T)$ is jointly convex. Consequently $C(T)$ is convex, and the optimal T may also be found using a golden section search. For this problem, the function $u(\tau)$ from Proposition 7 is $\hat{G}(R(\tau), L + \tau)$ where $\hat{G}(R, t) = hE[(R - X(t))^+]$. Since the cost function $C(T)$ retains the form $C(T) = \frac{K}{T} + \frac{1}{T} \int_0^T u(\tau) d\tau$, it is likely that sensitivity results similar to those in Section 4 might apply. However, a complete analysis for this situation would require a careful study of the (Q, r) policy under fill-rate constraints and of its deterministic counterpart⁵. Thus, sensitivity of the (R, T) policy under fill-rate constraints is a possible topic for future research.

7.2 Multi-Product / Multi-Echelon Problems

In this section, we apply our single-item, single stage results to analyze performance of (R, T) based heuristics for an N -item joint replenishment system and a special case of the single-item, two-echelon serial system.

7.2.1 The Joint Replenishment Problem (JRP)

In the JRP, we incur a major setup cost, K , every time *any* item is ordered and a minor setup cost, K_i , associated with each individual item, i , that is included in the order. This order/setup cost structure aims to coordinate order placement epochs for different items so as to save on fixed order costs relating to full-truckload transportation costs or quantity discounts for items ordered from the same supplier. Another application of this model involves coordination of the replacement of different spare parts, for instance, in the aircraft industry

⁵For instance, in this case, $u_d(\tau) = \tilde{H}\lambda\tau$ with $\tilde{H} = h\beta^2$. So the main sensitivity result in Theorem 11 would follow if $\lim_{\tau \rightarrow \infty} u'(\tau) \leq \tilde{H}\lambda$.

(Hopp & Kuo, [HK98]). In this case, the major setup cost corresponds to “grounding” the plane, and the minor setup costs correspond to specific items that are replaced. The presence of non-zero minor setup costs considerably complicates the problem since the optimal policy is not easily characterizable in terms of the set of items to order (or spare parts to replace) at different order epochs. Many heuristics, such as the (S, c, s) “can-order” policy (Federgruen *et al* [FGT84]) and the (Q, S_i) policy (Pantumsinchai [P92]), have been proposed for this complex stochastic problem. In the can-order policy, at a review epoch, if any item’s inventory position is at or below its reorder point s , then an order is placed to raise the inventory position of this item and all other items whose inventory position is at or below c to their corresponding order-up-to levels S . The (Q, S_i) policy triggers an order when the total amount demanded since the last replenishment order is Q , individual items are replenished up to S_i . Despite the existence of several such control policies for the JRP, no worst-case analysis exists for any of these heuristics. An exception is recent work by Eynan and Kropp [EK98] which addresses performance of a periodic review (R, T) system using a simplified cost function.

Heuristic for the JRP: Let T_L denote the “base planning period,” each item’s reorder interval will be a power-of-two multiple of T_L (where T_L itself is a decision variable). We allocate the joint setup cost, K , to the different items in such a way that item i ’s set-up cost equals $\alpha_i K + K_i$ for some non-negative α_i satisfying $\sum_i \alpha_i = 1$. Starting with $\alpha_i = 0$, the α_i are incremented sequentially as in [AI88] or [R86]. Only items with the smallest reorder intervals are allocated positive α_i . This is important because the joint setup cost portion $\alpha_i K$ will only be incurred when item i is ordered. If we want to correctly charge the full joint setup cost K when *any* item is ordered, we need to assign positive α to the items with the smallest reorder interval (highest ordering frequency). Given the allocation of K , the multi-item problem separates into N single-item problems. Let T_i^* be the optimal reorder interval for item i ’s (R, T) problem with fixed cost $\alpha_i K + K_i$. As in Roundy [R86], let $T_L = \min_i T_i^*$ and let k_i be the smallest non-negative integer satisfying $k_i \geq \log_2(T_i^*/T_L) - 0.5$. The reorder interval for item i is set to $T_i = 2^{k_i} T_L$ and the order-up-to level is $R_i = \mathcal{R}_i(T_i)$, where the latter denotes the optimal item i order-up-to level corresponding to reorder interval T_i .

Let C_{JRP}^* and C_{JRP}^H denote respectively, the optimal cost and the cost resulting from the above heuristic. Let $C_i^*(\alpha_i)$ denote the optimal cost of a (Q, r) policy for item i with setup cost $\alpha_i K + K_i$. Since, for any set $S \subseteq \{1, 2, \dots, N\}$ of ordered items, $K + \sum_{i \in S} K_i \geq$

$$\sum_{i \in S} (\alpha_i K + K_i): \quad C_{JRP}^* \stackrel{\text{by [AI88]}}{\geq} \sum_i C_i^*(\alpha_i) \quad (15)$$

Let $C_i(T_i)$ denote the long-run average cost incurred by item i , when its setup cost is $\alpha_i K + K_i$. Since $\frac{1}{\sqrt{2}} \leq \frac{T_i}{T_i^*} \leq \sqrt{2}$, by Theorems 11 and 15,

$$C_i(T_i) \leq \frac{1}{2} \left(\sqrt{2} + \frac{1}{\sqrt{2}} \right) C_i(T_i^*) = 1.06 C_i(T_i^*) \leq 1.06 \times \sqrt{2} C_i^*(\alpha_i) \quad (16)$$

Because we incrementally allocate the α_i in such a manner that only the items with the smallest reorder interval ($= \min_i T_i$) are assigned positive α_i values and since $\sum_i \alpha_i = 1$, every time an order is placed (i.e., every $\min_i T_i$ time units) our heuristic pays exactly K dollars for each major setup. Thus,

$$C_{JRP}^H = \sum_i C_i(T_i). \quad (17)$$

Proposition 16 *The worst-case performance of the above heuristic for the JRP is given by:*

$$C_{JRP}^H \stackrel{\text{by (15),(16),(17)}}{\leq} 1.06 \times \sqrt{2} C_{JRP}^* = 1.5 C_{JRP}^*.$$

Power-of-two based (R, T) policies may also be useful for a different single stage, multi-item production/inventory system in which resource constraints make coordination of order epochs for different items undesirable. For instance, there could be limited warehouse storage space, so that you prefer to stagger order delivery epochs. Building heuristics with worst-case performance guarantees for this problem environment with stochastic demand and different types of capacity constraints is another possible area for future research.

7.2.2 Two-Stage Serial System

Recent research on the single item, two-stage serial system has focused on developing effective heuristics based on the (Q, r) model. Under special circumstances (such as guaranteed availability or zero delivery lead times at the upstream stage), remarkable worst-case performance guarantees have been obtained by Atkins & De [AD93] and Chen [C94]. We apply the (R, T) policy to such two-stage systems. A typical inventory profile for such a system is shown in Figure 2. Clearly, the complicating factor is the interaction between echelons (e.g., the amount that can be delivered to the downstream stage is restricted by the availability of inventory at the upstream stage).

Let stage 2 be the upstream stage that interacts with a perfectly reliable supplier; and let stage 1 experience stationary stochastic customer demand at rate λ . The parameters at

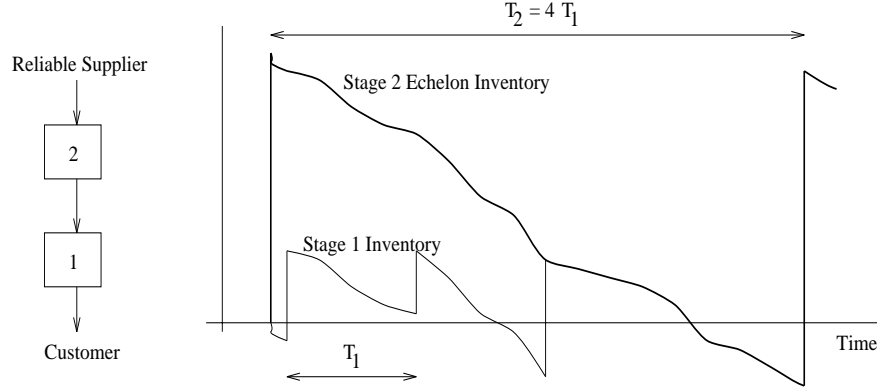


Figure 2: Typical Inventory Profile for Two-Stage Nested, (R, T) Policy

stages 1 and 2 are respectively: delivery (transportation) lead times, L_1, L_2 ; fixed order or setup cost, K_1, K_2 ; echelon holding cost rates, h_1, h_2 ; and backordering cost rate p_1 at stage 1. If $D(t_1, t_2)$ denotes cumulative demand over the interval $(t_1, t_2]$, the long-run average cost function for the two-stage (R, T) policy may be written as

$$C_{II}(R_2, T_2, R_1, T_1) = \frac{1}{T_2} \left\{ K_2 + \int_{L_2}^{L_2+T_2} h_2 E[(R_2 - D(0, t))^+] dt + \sum_{k=1}^n \left[K_1 + \int_{L_2+L_1+(k-1)T_1}^{L_2+L_1+kT_1} G(R_1^k, t) dt \right] \right\}, \quad (18)$$

where $T_2 = nT_1$ for some positive integer n , $R_1^k = \min(R_1, R_2 - D(0, L_2 + (k-1)T_1))$, and $G(R_1^k, t)$ is as defined in equation (4) with the appropriate holding and shortage cost parameters and $X(t)$ replaced by $D(L_2 + (k-1)T_1, t)$. Note that the first two terms in (18) correspond to costs incurred in one cycle of length T_2 at the upper stage; the terms within the summation correspond to costs incurred in the k^{th} cycle of length T_1 at stage 1. The term R_1^k incorporates the fact that, although stage 1 would ideally like to raise its inventory position up to R_1 , the actual increase in inventory position (as defined in Chen [C94] to be the on-hand – backlog + in-transit) may be smaller than R_1 when stage 2 has insufficient inventory.

Given T_1 and T_2 , the problem of minimizing $C_{II}(R_2, T_2, R_1, T_1)$ is similar to the discrete time, periodic review inventory problem with zero fixed costs, hence the cost function is convex in R_1, R_2 . Consequently, the optimal values of R_1 and R_2 denoted by $R_1(T_1, T_2)$ and $R_2(T_1, T_2)$ respectively, may be determined using gradient based search (where the gradients may be estimated using simulation). Let $C(T_1, T_2) \equiv C_{II}(R_2(T_1, T_2), T_2, R_1(T_1, T_2), T_1)$. The optimal values of (T_1, T_2) that minimize $C(T_1, T_2)$ may be found using a grid search (or a more efficient search when $C(T_1, T_2)$ is convex).

Analyzing the effectiveness of the above heuristic for the general case is complex (and needs to be carefully explored in a separate research paper). In this paper, we restrict attention to a special case with $K_2 = 0$. This may happen if the upper echelon has invested in information technology to facilitate the order placement process thereby reducing fixed order costs. (The lower echelon could be a small retailer who still uses an older order process thereby incurring positive fixed order costs.) This special case is the “reverse” of the case considered by Clark and Scarf [CS60] in the sense that they had $K_2 > 0$ with $K_1 = 0$. Whereas the optimal policy for the case in [CS60] is now well-known to be (Q, r) at stage 2 and order-up-to at stage 1, the optimal policy for our special case is not as well-known when $L_2 > 0$. The only simplification resulting from $K_2 = 0$ is that $T_2^* = T_1^*$ so that it is optimal for the two stages to follow the same cycle. In this case, the main problem is to determine, L_2 time units in advance, the amount of inventory stage 2 should carry so as to completely satisfy stage 1 orders whenever economically possible.

$K_2 = 0$: Since $T_2 = T_1$, the long-run average cost of the two-stage (R, T) policy may be written as $C(T) = C_{II}(R_1(T), T, R_2(T), T)$. Using arguments similar to Section 3, we note that this system behaves like a single stage system with a modified inventory cost consisting of holding costs at stage 2 and holding + shortage costs at stage 1. Using equation (18), it is possible to verify that $C(T) = \frac{1}{T}\{K_1 + \int_0^T G_{II}(R_1(\tau), \tau, R_2(\tau))d\tau\}$, where $G_{II}(R_1(\tau), \tau, R_2(\tau)) = h_2E[(R_2(\tau) - D(0, L_2 + \tau))^+] + G(R_1^1(\tau), L_1 + L_2 + \tau)$. Hence, analogous to Proposition 7, we can write $C(T) = \frac{1}{T}\{K_1 + \int_0^T u(\tau)d\tau\}$ with $u(\tau) = G_{II}(R_1(\tau), \tau, R_2(\tau))$ being increasing and convex in τ . Consequently, we can apply the appropriate two-stage version of Lemma 14 to show that

$$C(T^*) \leq \sqrt{2\lambda K_1 H_{12}} + \mathcal{G}(y_1^*, y_2^*), \quad (19)$$

where $H_{12} = \frac{(h_1+h_2)p_1}{h_1+h_2+p_1}$ and $\mathcal{G}(y_1^*, y_2^*) = G(R_1(0), 0, R_2(0))$. Using the inventory cost notation $\mathcal{G}_1(\cdot)$ and $\mathcal{G}_2(\cdot)$ in Chen and Zheng [CZ94], we see that their induced penalty lower bound on long-run average costs for this two-stage system is

$$\begin{aligned} C^{LB} &= \mathcal{C}_1(Q_1^*, r_1^*) + \mathcal{G}_2(y_2^*) \geq \sqrt{2\lambda K_1 H_{12}} + \mathcal{G}_1(y_1^*)^2 + \mathcal{G}_2(y_2^*) \\ &\geq \sqrt{2\lambda K_1 H_{12}} + \mathcal{G}(y_1^*, y_2^*)^2, \end{aligned} \quad (20)$$

where the first inequality follows from Lemma 13 and the last inequality follows by considering the case of $K_1 = K_2 = 0$ and observing that $\mathcal{G}(y_1^*, y_2^*) = \mathcal{G}_1(y_1^*) + \mathcal{G}_2(y_2^*)$.

From the upper and lower bounds on costs in (19) and (20), we conclude that the cost of our

(R, T) based heuristic is within $\sqrt{2}$ times the optimal, and no worse than $\frac{1}{2}(\sqrt{2} + \frac{1}{\sqrt{2}})\sqrt{2} = 1.5$ times the optimal when T is restricted to a power-of-two multiple of a base planning period. Thus, this problem can be effectively approximated using an (R, T) policy.

In summary, we see that for the JRP and the two-stage system considered above, the results from Section 3 may be used to provide provably effective (R, T) based policies. We believe that the effectiveness of the (R, T) policies will generalize to other systems. Our current research is investigating, computationally, the performance of power-of-two (R, T) -based heuristics for more general multi-item, multi-echelon inventory systems.

8 Conclusions

In this paper we have shown that the long-run average cost for the (R, T) inventory control policy is convex in the reorder interval T when $R = R(T)$, the optimal order-up-to-level corresponding to T . $R(T)$ is obtained by solving a newsvendor problem with an appropriately defined demand distribution; T^* may be found by an efficient golden section search. For continuous (normal) demand, the ratio of the cost of using $T = \alpha T^*$ instead of the optimal T^* is bounded by $\frac{1}{2}(\alpha + \frac{1}{\alpha})$, provided $R = R(T)$. Computational testing suggests that, in practice, use of the deterministic economic order interval, T_d^* , or a power-of-two approximation, instead of the optimal T^* does not increase costs significantly. As expected, the (R, T) policy is more expensive than the more responsive, (Q, r) policy. However, the difference in cost of the two policies may be bounded in the worst-case by 41.42% and is often on the order of a few percent. This implies that the (R, T) policy could be useful in situations where scheduling and coordination issues dominate. Based on the above results, we present a power-of-two, (R, T) -based heuristic with worst-case performance guarantees of 1.5 for the joint replenishment problem and a special case of two-stage serial systems.

A Appendix: Proofs

Proof of Propostion 3. From Section 6.A.1 in Shaked & Shanthikumar [SS94], we see that $X(t)$ is SICX iff $f_a(t) = E[\max(X(t) - a, 0)]$ is ICX in t for all a . For the Brownian motion, $X(t)$, with $\bar{\Phi}(z) = P(\text{Normal}(0, 1) \geq z)$, it is easy to verify that

$$f_a(t) = (\lambda t - a)\bar{\Phi}\left(\frac{a - \lambda t}{\sigma\sqrt{t}}\right) + \sigma\sqrt{t}\phi\left(\frac{a - \lambda t}{\sigma\sqrt{t}}\right),$$

and that the second derivative of $f_a(\cdot)$ w.r.t. t satisfies,

$$f_a''(t) = \frac{1}{4\sigma t^{5/2}} \phi\left(\frac{a - \lambda t}{\sigma\sqrt{t}}\right) [(a + \lambda t)^2 - \sigma^2 t] \geq 0,$$

where the non-negativity follows trivially when $a \geq 0$; when $a < 0$, $f_a''(t) = 0$, since⁶ $\phi\left(\frac{a - \lambda t}{\sigma\sqrt{t}}\right) \leq \phi(-3.5) = 0$. By Theorem 6.A.8 in [SS94], since $X(t)$ is SICX and $E[X(t)] = \lambda t$, $X(t)$ is SIL. ■

Proof of Lemma 5. (i) is trivially true. (ii) Since, both $X(t)$ and $\mathcal{T}(T) = L + \mathcal{U}(0, 1)T$ are SIL, by preservation Theorem 6.A.13 in Shaked and Shanthikumar [SS94], $Y(T) = X(\mathcal{T}(T))$ is SIL in T . (iii) Similarly, it is easy to see that, if $T_x \geq T_y$, then $\tilde{Y}(\alpha)$ is SIL, otherwise $\tilde{Y}(\alpha)$ is SDL in α . By Proposition 3.7 in [SS88], we see that if $\tilde{Y}(\alpha)$ is SIL or SDL, then it is SCX in α . (iv) As in (iii), $\tilde{Y}(\alpha) - \tilde{R}(\alpha)$ is either SIL or SDL in α , hence SCX. ■

Proof of Proposition 9. If $U(T)$ is convex, then $u'(\tau) = U''(\tau) \geq 0$ and $u(\tau)$ is non-decreasing. Convexity of $U(T)$ may be proven as follows: $U(T) = TH(R(T), T)$, so that $U''(T) = TH''(R(T), T) + 2H'(R(T), T)$. By convexity of $\min_R H(R, T)$, $H''(R(T), T) \geq 0$. Further, since $H(R(T), T)$ denotes optimal total inventory costs incurred over a cycle of length T , $H'(R(T), T) \geq 0$ (see Lemma 17 below). Consequently, $U''(T) \geq 0$.

To prove convexity of $u(\tau)$, note that $X(L+T) = Y(T) + Z(T)$ where $Z(T) \sim X(\mathcal{U}(0, 1)T)$ is SICX. Further, $u(\tau) = G(R(\tau), L + \tau) \equiv E[g(Y(\tau) + Z(\tau) - R(\tau))]$, where $g(x) = [h(-x)^+ + p(x)^+]$. Since $E[g(Y(\tau) - R(\tau))] = H(R(\tau), \tau)$ is convex, non-decreasing in τ , and $Z(\tau)$ is SICX, extending results from [SS88] yields that $E[g((Y(\tau) - R(\tau)) + Z(\tau))] = u(\tau)$ is convex in τ . (Note: For the Brownian motion demand model, Proposition 9 is proven directly in Appendix B, without use of Stochastic Convexity.) ■

Lemma 17 $TH'(R(T), T) = (h + p)E_{Z(T)} \left[\int_{R(T)-Z(T)}^{R(T)} (y - R(T) + Z(T)) \gamma(y, T) dy \right] \geq 0$.

Proof. For simplicity of presentation we replace $H(R(T), T)$ with $H(\cdot)$. By definition of $H(\cdot)$ and the fact that $R(T)$ minimizes $H(R, T)$, we conclude that

$$\begin{aligned} H'(\cdot) = \frac{\partial}{\partial T} H(\cdot) &= -\frac{h\lambda}{2} + (h + p) \int_{y \geq R(T)} (y - R(T)) \frac{d}{dT} \gamma(y, T) dy \\ &= -\frac{h\lambda}{2} + \frac{(h + p)}{T} \left[E[(X(L + T) - R(T))^+] - E[(Y(T) - R(T))^+] \right], \end{aligned} \quad (21)$$

⁶This may be directly verified by noting that for $a < 0$ with $X(t) \geq 0$, $E[\max(X(t) - a, 0)] = E[X(t) + |a|] = \lambda t + |a|$ is linear in t .

where the second equality follows from the definition of $\gamma(y, T)$. Since $X(L + T) \stackrel{d}{=} Y(T) + Z(T)$, the term within the square brackets in equation (21) above is

$$\begin{aligned} & E[(Y(T) - R(T) + Z(T))^+] - E[(Y(T) - R(T))^+] \\ &= E_{Z(T)} \left[\int_{y \geq R(T) - Z(T)} (y - R(T) + Z(T)) \gamma(y, T) dy \right] - \int_{y \geq R(T)} (y - R(T)) \gamma(y, T) dy \\ &= E_{Z(T)} \left[\int_{R(T) - Z(T)}^{R(T)} (y - R(T) + Z(T)) \gamma(y, T) dy + \int_{y \geq R(T)} Z(T) \gamma(y, T) dy \right]. \end{aligned} \quad (22)$$

Now the required result follows from (21) and (22), using $\int_{y \geq R(T)} \gamma(y, T) dy = P(Y(T) \geq R(T)) = \frac{h}{h+p}$ and $E[Z(T)] = \frac{\lambda T}{2}$. \blacksquare

Proof of Proposition 10. By Proposition 7 and equation (6),

$$u(\tau) = G(R(\tau), L + \tau) = p(\lambda(L + \tau) - R(\tau)) + (h + p) \int_{x \leq R(\tau)} \Psi(x, L + \tau) dx.$$

For the Brownian motion case, it is easy to show that, for $w_\tau = \frac{R(\tau) - \lambda(L + \tau)}{\sigma \sqrt{L + \tau}}$,

$$u(\tau) = h(R(\tau) - \lambda(L + \tau)) + (h + p) \left[(\lambda(L + \tau) - R(\tau)) \bar{\Phi}(w_\tau) + \sigma \sqrt{L + \tau} \phi(w_\tau) \right].$$

Differentiating, $u'(\tau) = h(R'(\tau) - \lambda) + (h + p) \left[(\lambda - R'(\tau)) \bar{\Phi}(w_\tau) + \frac{1}{2} \frac{\sigma}{\sqrt{L + \tau}} \phi(w_\tau) \right]$, and

$$\begin{aligned} \lim_{\tau \rightarrow \infty} u'(\tau) &= \lim_{\tau \rightarrow \infty} (\lambda - R'(\tau)) \left[(h + p) \bar{\Phi}(w_\tau) - h \right] + 0, \\ &\stackrel{(11)}{=} \lim_{\tau \rightarrow \infty} \frac{\lambda h}{(h + p) \bar{\Phi}(w_\tau)} \left[(h + p) \bar{\Phi}(w_\tau) - h \right], \\ &= \frac{\lambda h}{(h + p)} \left[(h + p) - h / \bar{\Phi}(w_\tau) \right] \leq \frac{\lambda h}{h + p} [(h + p) - h]. \end{aligned}$$

Thus, $\lim_{\tau \rightarrow \infty} u'(\tau) \leq \frac{hp\lambda}{h+p} = H\lambda$. By convexity of $u(T)$, $u'(T) \leq \lim_{\tau \rightarrow \infty} u'(\tau)$, which yields the required result. \blacksquare

Lemma 18 For any $\alpha > 0$, $T > 0$, $\int_T^{\alpha T} u(\tau) d\tau \leq \frac{1}{2}(\alpha^2 - 1) T u(T)$.

Proof. See proof of Lemma 9 in Zheng [Z92]. \blacksquare

B Appendix: Normal Demand Model

Wlog, we scale time so that $T \geq 1$. Let $\bar{\Phi}(x)$ denote the complementary CDF for the standard normal. Let $v \equiv v(T) = R(T) - \lambda(L + T)$ and $w \equiv w(T) = \frac{v(T)}{\sigma \sqrt{L + T}}$. Note that $v' = R'(T) - \lambda$, and $w' = \frac{1}{\sigma \sqrt{L + T}} (R'(T) - \frac{\lambda}{2} - \frac{R(T)}{2(L + T)})$.

Proposition 19 For normal demand, $u(T)$ is non-decreasing and convex.

Proof. For normal demand,

$$\begin{aligned} u(T) &\equiv h(R(T) - \lambda(L + T)) + (h + p) \int_{R(T)}^{\infty} (x - R(T))\psi(x, L + T)dx \\ &= hv + (h + p)[\sigma\sqrt{L + T}\phi(w) - v\bar{\Phi}(w)] \quad (\text{by \# 25., Appendix 4, [HW63]}). \end{aligned}$$

$$\text{Hence } u'(T) = [h - (h + p)\bar{\Phi}(w)]v' + (h + p)\frac{\sigma}{2\sqrt{L + T}}\phi(w). \quad (23)$$

Since $u'(0) \geq 0$, $u''(T) \geq 0$ will imply that $u'(T) \geq 0$ for all $T \geq 0$ and $u(T)$ will be non-decreasing (and convex). Clearly, $u''(T)$ depends on $R''(T)$. By equation(9), $R'(T) \int_0^T \psi(R(T), L + t)dt = \bar{\Phi}(w) - \frac{h}{h+p}$. Differentiating,

$$R''(T) \int_0^T \psi(R(T), L + t)dt + R'(T) \frac{d}{dT} \int_0^T \psi(R(T), L + t)dt = -\phi(w)w'. \quad (24)$$

Since $\psi(R(T), t) = \frac{1}{\sigma\sqrt{L+T}}\phi(\frac{R-\lambda t}{\sigma\sqrt{t}})$, from # 9., Appendix 4, [HW63],

$$\int_0^T \psi(R, L+t)dt = \frac{1}{\lambda}[W(R, L+T) - W(R, L)] \quad \text{where } W(R, t) = \bar{\Phi}\left(\frac{R - \lambda t}{\sigma\sqrt{t}}\right) - e^{2\lambda R/\sigma^2} \bar{\Phi}\left(\frac{R + \lambda t}{\sigma\sqrt{t}}\right).$$

By Leibnitz rule, $\frac{d}{dT} \int_0^T \psi(R(T), L + t)dt = \psi(R(T), L + T) + \frac{R'(T)}{\lambda} \frac{d}{dR}[W(R, L + T) - W(R, L)]$. Given our Brownian demand model, it is easy to verify that $\frac{d}{dR}[W(R, L + T) - W(R, L)] \geq 0$. Consequently, $\frac{d}{dT} \int_0^T \psi(R(T), L + t)dt \geq \psi(R(T), L + T)$. Substituting this into equation (24), we get

$$R''(T) \int_0^T \psi(R(T), L + t)dt \leq -\phi(w)w' - R'(T)\psi(R(T), L + T). \quad (25)$$

In summary, we see that:

$$R''(T) \frac{[\bar{\Phi}(w) - \frac{h}{h+p}]}{R'(T)} \stackrel{(9)}{=} R''(T) \int_L^{L+T} \psi(R(T), t)dt \stackrel{(25)}{\leq} -\phi(w)w' - R'(T) \frac{\phi(w)}{\sigma\sqrt{L + T}}. \quad (26)$$

Differentiating equation (23) and substituting $v''(T) = R''(T)$, we get

$$\begin{aligned} \frac{u''(T)}{h+p} &= -R''(T)[\bar{\Phi}(w) - \frac{h}{h+p}] + [\phi(w)w']v' - \frac{\sigma}{2\sqrt{L+T}}w\phi(w)w' - \frac{\sigma}{4(L+T)^{3/2}}\phi(w), \\ &\stackrel{(26)}{\geq} R'(T)[\phi(w)w' + R'(T)\frac{\phi(w)}{\sigma\sqrt{L+T}}] + \phi(w)w'[v' - \frac{\sigma}{2\sqrt{L+T}}w] - \frac{\sigma}{4(L+T)^{3/2}}\phi(w). \end{aligned}$$

It is easy to verify that $[v' - \frac{\sigma}{2\sqrt{L+T}}w] = \sigma\sqrt{L + T}w'$. Thus,

$$\begin{aligned} \frac{u''(T)}{h+p} &= \frac{1}{\sigma\sqrt{L+T}}\phi(w)[\sigma\sqrt{L + T}R'(T)w' + R'(T)^2 + [\sigma\sqrt{L + T}w']^2 - \frac{\sigma^2}{4(L+T)}] \\ &= \frac{1}{\sigma\sqrt{L+T}}\phi(w)[(\sigma\sqrt{L + T}w' + \frac{R'(T)}{2})^2 + \frac{3}{4}R'(T)^2 - \frac{\sigma^2}{4(L+T)}] \geq 0. \end{aligned}$$

The non-negativity above follows by considering two cases: (i) $3R'(T)^2 \geq \frac{\sigma^2}{(L+T)}$, for which

the right hand side is clearly non-negative, and, (ii) $3R'(T)^2 < \frac{\sigma^2}{(L+T)}$, in which case, by the equation for w' (beginning of this appendix), we see that $(\sigma\sqrt{L+T}w' + \frac{R'(T)}{2})^2 \geq \frac{\sigma^2}{4(L+T)}$. ■

Acknowledgements. We thank Karla Bourland, Awi Federgruen, Mike Fry, Jack Muckstadt, and Alan Scheller-Wolf for constructive comments on a previous version of this paper. We thank M&SOM deputy editor, Paul Zipkin, the senior editor, and two anonymous referees for their help in improving this paper.

References

- [AD93] Atkins, D. and S. De, “94%-Effective Lot-Sizing for a Two Stage Serial Production / Inventory System with Stochastic Demand,” Working Paper 92-MSC-001, Faculty of Commerce, UBC, Canada, 1992.
- [AI88] Atkins, D. and P. Iyogun, “Periodic Versus ‘Can-Order’ Policies for Coordinated Multi-Item Inventory Systems,” *Mgmt. Sci.* 34, 791-796, 1988.
- [BSS93] Bazaraa, M. S., H. D. Sherali, and C. M. Shetty, “*Nonlinear Programming: Theory and Algorithms*”, 2nd Edition, John Wiley & Sons, NY, 1993.
- [C01] Cachon, G. “Managing a retailer’s shelf space, inventory, and transportation,” *Mfg. and Serv. Opns. Mgmt.*, 3, 3, pp. 211-29, 2001.
- [CSY94] Chang, C.-S., J.G. Shanthikumar, and D.D. Yao, “Stochastic Convexity and Stochastic Majorization,” Chapter 5 in “*Stochastic Modeling and Analysis of Manufacturing Systems*,” Ed. D.D. Yao, Springer-Verlag, 1994.
- [C94] Chen, F., “94%-Effective Policies in a Two-Stage Serial System with Stochastic Demand,” *Mgmt. Sci.* 45, 1679-1696.
- [CZ94] Chen, F., and Y-S. Zheng, “Lower Bounds for Multi-echelon Stochastic Inventory Systems,” *Mgmt. Sci.*, 40, 11, pp. 1426-1443, 1994.
- [CS60] Clark, A.J., and H. Scarf, “Optimal Policies for a Multi-Echelon Inventory Problem,” *Mgmt. Sci.*, 6, 475-490, 1960.
- [EK98] Eynan A., and D. Kropp, “Periodic Review and Joint Replenishment in Stochastic Demand Environments,” *IIE Trans.*, 30, 11, 1025-1033, 1998.
- [FGT84] Federgruen A., H. Groenvelt, and H. C. Tijms, “Coordinated Replenishments in a Multi-Item Inventory System with Compound Poisson Demands,” *Mgmt. Sci.*, 30, 3, pp. 344-357, 1984.

- [F71] Feller, W., “*An Introduction to Probability Theory and its Applications*,” Volume II, 2nd Edition, John Wiley & Sons, 1971.
- [G98] Gallego, G., “New Bounds on (Q, r) Policies,” *Mgmt. Sci.*, 44, 2, 219-233, 1998.
- [G96] Graves, S.C., “A Multi-Echelon Inventory Model with Fixed Replenishment Intervals,” *Mgmt. Sci.*, 42, 1, pp. 1-18, 1996.
- [GW00] Graves, S.C., and S.P. Willems, “Optimizing strategic safety stock placement in supply chains,” *Mfg. and Serv. Opns. Mgmt.*, 2, 1, pp. 68-83, 2000.
- [HK98] Hopp, W. J. and Y-L. Kuo, “Heuristics for Multi-Component Joint Replacement: Applications to Aircraft Engine Maintenance,” *Nav. Res. Log.*, 45, 5, 45-58, 1998.
- [HW63] Hadley, G. and T. M. Whitin, “*Analysis of Inventory Systems*,” Prentice Hall, Englewood Cliffs, NJ, 1963.
- [HS84] Heyman, D.P., and M.J. Sobel, “*Stochastic Models in OR*,” Vol II, McGraw-Hill, NY, 1984.
- [MS00] Moses, M., and S. Seshadri, “Policy mechanisms for supply chain coordination,” *IIE Trans.* 32, 3, pp. 245-62, 2000.
- [MR93] Muckstadt, J.A., and R.O. Roundy, “Analysis of Multistage Production Systems,” Chapter 2 in *Logistics of Production and Inventory*, Handbooks of OR & MS, North-Holland, 1993.
- [P92] Pantumsinchai, P., “A Comparison of Three Joint Ordering Inventory Policies,” *Dec. Sci.*, 23, pp. 111-127, 1992.
- [RST00] Rao, U.S., A. Scheller-Wolf, and S. Tayur, “Development of a Rapid-Response Supply Chain at Caterpillar,” *Opns. Res.*, 48, 189-204, 2000.
- [R70] Rockafellar, R. T., “*Convex Analysis*,” Princeton University Press, Princeton, NJ, 1970.
- [R86] Roundy, R., “A 98% Effective Lot Sizing Rule for Multi-Product, Multi-Stage Production Inventory Systems,” *Math. of OR* 11, 699,727, 1986.
- [SS78] Serfozo, R. and S. Stidham, “Semi-Stationary Clearing Processes,” *Stoch. Proc. Appl.*, 6, 165-178, 1978.
- [SS88] Shaked, M. and G. Shanthikumar, “Stochastic Convexity and its Applications,” *Adv. Appl. Prob.* 20, 427-488, 1988.

- [SS94] Shaked, M. and G. Shanthikumar, "Stochastic Orders and Their Applications," Academic Press, Inc., NY, 1994.
- [SPP98] Silver, E., D.F. Pyke, and R. Peterson, "*Inventory Management and Production Planning and Scheduling*", 3rd Edition, John Wiley & Sons, NY, 1998.
- [TV01] Teunter, R., and D. Vlachos, "A periodic review inventory system with emergency replenishments," *Mgmt. Sci.*, 47, 3, pp. 415-29, 2001.
- [Z98] Zhang, Hongtao, "A Note on the Convexity of Service-Level Measures of the (r, q) System," *Mgmt. Sci.* 44, 431-432, 1998.
- [Z92] Zheng, Y-S., "Properties of Stochastic Inventory Systems," *Mgmt. Sci.* 38, 87-103, 1992.
- [ZC92] Zheng, Y-S., and F. Chen, "Inventory Policies with Quantized Ordering," *Nav. Res. Log.*, 39, pp. 285-305, 1992.
- [ZF91] Zheng, Y-S, and A. Federgruen, "Finding optimal (s, S) policies is about as simple as evaluating a single policy," *Opns. Res.* 39, 654-665, 1991.
- [Z86] Zipkin, P., "Inventory Service Level Measures: Convexity and Approximation," *Mgmt. Sci.* 32, 975-981, 1986.
- [Z00] Zipkin, P., "Foundations of Inventory Management", McGraw-Hill, 2000.