RAIRO Operations Research
RAIRO Oper. Res. **35** (2001) 37-70

OPTIMAL AND NEAR-OPTIMAL (s, S) INVENTORY POLICIES FOR LEVY DEMAND PROCESSES *, **

Robin O. Roundy[1] and Gennady Samorodnitsky[1]

Communicated by N. Kaio

Abstract. *A Levy jump process* is a continuous-time, real-valued stochastic process which has independent and stationary increments, with no Brownian component. We study some of the fundamental properties of Levy jump processes and develop (s, S) inventory models for them. Of particular interest to us is the gamma-distributed Levy process, in which the demand that occurs in a fixed period of time has a gamma distribution. We study the relevant properties of these processes, and we develop a quadratically convergent algorithm for finding optimal (s, S) policies. We develop a simpler heuristic policy and derive a bound on its relative cost. For the gamma-distributed Levy process this bound is 7.9% if backordering unfilled demand is at least twice as expensive as holding inventory. Most easily-computed (s, S) inventory policies assume the inventory position to be uniform and assume that there is no overshoot. Our tests indicate that these assumptions are dangerous when the coefficient of variation of the demand that occurs in the reorder interval is less than one. This is often the case for low-demand parts that experience sporadic or spiky demand. As long as the coefficient of variation of the demand that occurs in one reorder interval is at least one, and the service level is reasonably high, all of the polices we tested work very well. However even in this region it is often the case that the standard Hadley–Whitin cost function fails to have a local minimum.

Received May, 2000. Accepted February, 2001.

* *Research supported by the NSF grant DDM-921704.*

** *Research supported by the NSF grant DMS-94-00535 and the NSA grant MDA904-95-H-1036.*

[1] School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853, U.S.A.

1. Introduction

We consider (s,S) inventory models in which both time and inventory are modelled as continuous quantities, lead times are deterministic, and customer service is modelled through cost minimization rather than constraints on service levels.

We begin be discussing models in which the inventory position has a uniform distribution. The most popular model for (s,S) and (R,Q) policies is originally due to Hadley and Whitin [18]. This model has been the mainstay of introductory textbooks on inventory theory for over 20 years (see, for example, Johnson and Montgomery [20], Nahmias [21] and Vollman *et al.* [30]). This model assumes that the inventory position lands on the reorder point rather then jumping over it, so the distinction between (s,S) policies and (R,Q) policies disappears. This model also assumes that the inventory position is uniformly distributed between s and S. The cost function makes use of approximations (see Sect. 4). This model is conceptually and computationally simple. Any demand distribution can be used, continuous or discrete. The computations required are relatively simple, and the model gives good solutions for most real-world inventory systems.

The main disadvantage of the Hadley–Whitin model is its robustness. When the backorder cost is sufficiently low or the order cost is sufficiently high the cost function can fail to have a local minimum. For example, suppose that the holding cost is 1 dollar per item per day, the mean demand rate is 1 item per day, the order cost is K dollars, the backorder cost is $\hat{p}$ dollars per item, and the demand that occurs during one lead time is uniformly distributed over the interval $[0,2]$. If $\hat{p}(\hat{p}-2) \leq 2K$ the cost function has no local minimum ((31) in Sect. 4 is non-decreasing). Thus if $\hat{p}=10$ and $K \geq 40$, or if $\hat{p}=40$ and $K \geq 760$, the cost function has no local minimum.

More recently Zheng developed an elegant (R,Q) inventory model [32]. This model differs from that of Hadley and Whitin in that Zheng uses a time-weighted backorder cost of p dollars per item per day rather than $\hat{p}$ dollars per item. Zheng assumes that the cumulative demand process is non-decreasing, has continuous sample paths, and has identically distributed increments. The continuity of the sample paths implies that the inventory position lands on the reorder point rather than jumping over it, and that the inventory position is uniformly distributed between s and S. As in the Hadley–Whitin model, the distinction between (s,S) policies and (R,Q) policies disappears.

Zheng's model is gaining in popularity. It makes no approximations in the cost function, so it is not surprising that Zheng's model is more robust than the Hadley–Whitin model. It is more complex than the Hadley–Whitin model, both conceptually and computationally, but it is simple enough to teach in introductory courses and efficient enough to use in large, real-world inventory systems. Zheng's algorithms work with any demand distribution.

Chen and Zheng [5] study (R,Q) policies for a system in which either time is discrete or the demand process is compound Poisson. For these systems (R,Q) policies are not optimal, and the inventory position often jumps over the reorder point rather than landing on it. However the inventory position has a uniform distribution. The authors give an efficient algorithm for computing an (R, Q) policy.

A *Levy jump process* is a continuous-time, real-valued stochastic process which has independent and identically distributed (*i.i.d.*) increments, with no Brownian component. According to the Levy Decomposition theorem (Hida [19], p. 45), any real-valued Levy process with *i.i.d.* increments can be expressed as $X(t) = \mu \cdot t + \sigma \cdot B(t) + J(t)$, where $\mu \cdot t$ is a deterministic drift, $B(t)$ is a Brownian motion and $J(t)$ is a Levy jump process. Zheng's cumulative demand process is assumed to be non-decreasing and continuous. Brownian motion is not non-decreasing and Levy jump processes are discontinuous, so according to this theorem either Zheng's demand process is deterministic or it fails to have independent increments.

Two different demand processes with dependent, identically-distributed increments that satisfy Zheng's assumptions have been studied (see Browne and Zipkin [4] and Serfozo and Stidham [27]). For these processes a stationary (s, S) policy leads to a uniform distributuion for the inventory position. Browne and Zipkin [4] give an algorithm for computing optimal (s, S) policies with phase-type demands. However (s, S) policies are no longer an optimal class because past demands are correlated with future demands.

We now turn to models in which the distribution of the inventory position is not uniform. Beyer and Sethi [2] and Sahin [26] make theoretical properties, but do not give algorithms. When the demand process is either a discrete-time process or a compound Poisson process, there are efficient policy-improvement algorithms that compute optimal (s, S) policies (see Federgruen and Zipkin [9,10] and Zheng and Federgruen [31]). These authors use clever, efficient versions of the policy improvement algorithm. In each iteration a renewal-type equation defines the inventory position. These equations can be solved directly when demand has either a discrete distribution or a continuous phase-type distribution. Fu [12] estimates the derivative of the cost with respect to s and S, enabling standard optimization algorithms. Hu [13] computes a differnt derivative with respect to s and S. With the compound Poisson processes there is a positive probability that the total demand incurred in any time interval of fixed length is equal to zero.

There are a numerous related models and results. See, for example, Axsater [1], Cheung [6], Federgruen and Zipkin [8], Gallego [14]. For recent work on service-constrained models see Robinson [24] and Gallego and Boyaci [15–17]. Song and Zipkin [29] and Bollapragada [3] study dynamic (s, S) models for systems with dependent demand.

We study (s, S) inventory models with Levy jump demand processes. Levy jump processes are a rich class of stochastic processes in which the inventory position usually jumps over the reorder point rather than landing exactly on it. Compound Poisson processes are Levy jump processes in which there is a positive probability of zero demand in a time interval of unit length. As far as we know, this is the first paper to study (s, S) inventory models for Levy jump demands that are not compound Poisson.

We prove some fundamental properties of (s, S) inventory models with Levy jump demands. We develop a quadratically convergent algorithm for finding optimal (s, S) policies for Levy jump processes. We develop a simpler policy called the Mass Uniform heuristic, and derive a bound on its relative cost. We do computational tests for the gamma-distributed Levy process, in which the demand that occurs in a fixed period of time has a gamma distribution. For the gamma-distributed Levy process this bound is 7.9% if backordering unfilled demand is at least twice as expensive as holding inventory. Our quadratically-convergent algorithm for computing optimal policies can also be used to compute Zheng's (R, Q) policy.

Another of our goals in initiating this research was to assess the robustness of Zheng's algorithm, by testing its performance for demand processes that overshoot the reorder point and have non-uniform distributions for the inventory position. Our computational tests indicate that it is very robust, much more so than the cost-minimization versions of the Hadley–Whitin model. However it can produce poor policies in some realistic scenarios.

Overview

This paper is organized as follows. In Section 2 we develop the key properties of general Levy demand processes and, specifically, of the gamma-distributed Levy process. In Section 3 we develop our inventory model for (s, S) policies for Levy demand processes. We also present a quadratically-convergent algorithm for computing optimal policies, and we discuss service levels. In Section 4 we develop the Mass Uniform heuristic for computing (s, S) policies, we prove that the relative cost of the Mass Uniform heuristic is at most 7.9% if backordering unfilled demand is at least twice as expensive as holding inventory. We develop bounds on the performance of the (s, S) policy derived from Zheng's (R, Q) policy, and we briefly describe the other policies that we have tested. In Section 5 we summarize our computational experiments, and in Section 6 we draw our conclusions. Appendix 6 contains a glossary of notation. Most proofs omitted from this paper, but they are found in [25].

2. Levy demand processes

We begin this section a definition.

Definition 1. A *Levy demand process* $D(t)$ is a non-decreasing right-continuous Levy process with no Brownian component, satisfying $D(0) = 0$ and $D(t) < \infty$ for all t.

We use Levy demand processes to model cumulative demand. Thus $D(t)$ is the total demand that occurs in the time interval $[0, t]$ and $D(t, u)$ is the total demand that occurs in the time interval $(t, u]$. By selecting the units of measure in our inventory models appropriately, we assume without loss of generality that

$$E[D(t)] = t.$$

Examples of Levy demand processes include the compound Poisson processes. In a Levy demand process, demand for randomly-sized quantities of inventory occurs instantaneously at random points in time, creating "demands" or "jumps" (discontinuities in the cumulative demand process). If the demand quantities are all integer multiples of some number, we say that the jump sizes are arithmetic (Feller [11], p. 360). Some of our results require non-arithmetic jump sizes.

Property 1. The jump sizes are non-arithmetic.

Any of three functions can be used to characterize a Levy demand process $D(t)$. The first of these comes from observing the time epochs at which a demand of size greater than x occurs. Since the demand process $D(t)$ has i.i.d. increments, these time epochs form a Poisson process. $\psi^+(x)$ is the rate at which demands of size greater than x occur. Let $F_{D(t)}(x)$ be the distribution function of $D(t)$. Algebraically we have

$$\psi^+(x) = \lim_{t\to 0} \frac{1}{t}\left[1 - F_{D(t)}(x)\right]. \tag{1}$$

(See Feller, p. 302, Th. 1. Note that the distributions of $D(t)$ for $t > 0$ form a convolution semi-group as defined in Feller p. 293.) Clearly $\psi^+(x)$ is non-increasing. The mean rate per unit time at which demand occurs is $\int_0^\infty x\,\mathrm{d}\{-\psi^+(x)\} = E[D(1)] = 1$. Because $D(t)$ is well-defined and non-decreasing, $\mathrm{limsup}_{x\to 0} x\,\psi^+(x) = 0$. If $\mathrm{limsup}_{x\to\infty} x\,\psi^+(x) > 0$ we have $E[D(1)] = \infty$, which is contrary to our assumptions. Integration by parts leads to

$$1 = E[D(1)] = \int_0^\infty x\,\mathrm{d}\{-\psi^+(x)\} = \int_0^\infty \psi^+(x)\,\mathrm{d}x. \tag{2}$$

A Levy demand process can be constructed from any non-negative, non-increasing function $\psi^+(x)$ satisfying (2). Compound Poisson processes are the Levy demand processes in which demands of any size occur at a finite rate, *i.e.*, $\psi^+(0) < \infty$. For a compound Poisson process $P(D(t) = 0) > 0$ for all $t \geq 0$ and $\psi^+(x)/\psi^+(0)$ is the probability that the size of a given jump is greater than x. On the other hand, if $\psi^+(x) \to \infty$ as $x \to 0$ then any open interval on the time axis contains a countably infinite number of jumps, and $P(D(t) = 0) = 0$ for all $t > 0$.

The measure $\mathrm{d}\{-\psi^+(x)\}$ is called the Levy measure. We define the random variables J and V by

$$F_J(z) = \int_0^z x\,\mathrm{d}\{-\psi^+(x)\} \quad \text{and} \quad F_V(z) = \int_0^z \psi^+(x)\,\mathrm{d}x. \tag{3}$$

Let $1_S = 1$ if S is true and $1_S = 0$ otherwise. The following lemma allows us to interpret J as the weighted jump size.

Lemma 1. *For all* $T > 0, E\left(\frac{1}{T}\int_0^T 1_{(D(t)-D(t^-)>z)}\,\mathrm{d}\{D(t)\}\right) = \bar{F}_J(z)$.

All proofs are found in [25]. To illustrate Lemma 1, suppose that demands of size 1 arrive at rate 1/3, and that demands of size 2 arrive at rate 1/3. Then $\psi^+(x) = 2/3$ for $0 \le x < 1$, $\psi^+(x) = 1/3$ for $1 \le x < 2$, and $\psi^+(x) = 0$ otherwise. Half of the jumps are of size 1 and half are of size 2, but J weights the probabilities of the jumps by their size, so $J = 1$ with probability 1/3, and $J = 2$ with probability 2/3.

The second function that is used to characterize Levy demand processes is the distribution of $D(t)$. Because $D(t)$ has stationary, independent increments, the Laplace transform $\mathcal{L}_{D(t)}(\gamma)$ of $D(t)$ satisfies

$$\mathcal{L}_{D(t)}(\gamma) = [\mathcal{L}_{D(1)}(\gamma)]^t. \tag{4}$$

Since $D(t)$ is a non-trivial, non-negative random variable,

$$\mathcal{L}_{D(t)}(\gamma) < 1 \text{ for } t > 0. \tag{5}$$

By Hida [19][2],

$$\mathcal{L}_{D(t)}(\gamma) = \mathrm{e}^{-t\int_0^\infty \gamma\,\mathrm{e}^{-\gamma x}\psi^+(x)\,\mathrm{d}x} = \mathrm{e}^{-t\int_0^\infty (1-\mathrm{e}^{-\gamma x})\,\mathrm{d}\{-\psi^+(x)\}}. \tag{6}$$

The third function that can be used to characterize Levy demand processes is the expected length of time $\theta(x) \equiv E[D^{-1}(x)]$ required to accumulate x units of demand, starting from a given point in time. $\theta(x)$ is related to the steady-state distribution of the inventory position. Clearly $\theta(0) = 0$. $\theta(0^+) > 0$ if and only if $\psi^+(0) < \infty$, *i.e.*, if and only in $D(t)$ is a compound Poisson process. $\theta(x)$ is left-continuous, and

$$\theta(x)/x \to 1 \quad \text{as} \quad x \to \infty. \tag{7}$$

One might be tempted to conjecture that $\theta(x) \equiv x$, but this is typically not the case.

Lemma 2. $\theta(x)$ *satisfies* $\theta(x) = E[D(D^{-1}(x))] \ge x$ *for all* $x > 0$. *If Property 1 holds then* $\theta(x) - x \to E[V]$ *as* $x \to \infty$.

In general $E[V]$ can be either finite or infinite. Lemma 2 holds in either case. Note that $D(D^{-1}(x)) \ge x$ by definition.

Intuitively, the fact that $\theta(x) > x$ can be explained as follows. If we are told that $D(\tau) = d$ then we know that $D(t)$ lands on d rather than jumping over d. This fact increases the expected length of time that $D(t)$ spends in the interval $[d, d+\delta)$, $\delta > 0$.

[2]Theorem 3.3 (p. 42). Hida's $\mathrm{d}\{n(u)\}$ is our $\mathrm{d}\{-\psi^+(u)\}$. His $X(t,\omega) + \int_0^\infty \frac{t\,u}{1+u^2}\,\mathrm{d}\{n(u)\}$ is our $D(t)$.

The overshoot $D(D^{-1}(x)) - x$ is the amount by which a Levy demand process over-shoots a given value x. The expectation of the overshoot is $\theta(x) - x$. Note that many inventory models, including those of Zheng and Hadley and Whitin, assume that there is no overshoot, so $\theta(x) - x$ is equal to zero for all x. Lemma 2 describes the expectation of the overshoot. The following lemma allows us to interpret V as the asymptotic distribution of the overshoot.

Lemma 3. *If Property 1 holds then* $P(D(D^{-1}(x)) - x > u) \to \int_u^\infty \psi^+(z)\,\mathrm{d}z = \bar{F}_V(u)$ *as* $x \to \infty$.

Equation (3) and the integration by parts formula imply that $F_V(z) = F_J(z) + z\cdot\psi^+(z) > F_J(z)$ and $E[J] = 2\cdot E[V]$. The means are finite if $\int_0^\infty x\,\psi^+(x)\,\mathrm{d}x < \infty$. Thus the asymptotic overshoot V is stochastically smaller than the weighted jump size J, and has a mean that is half of the mean of the weighted jump size.

The Laplace transform of $\theta(x)$ is

$$\begin{aligned}
\mathcal{L}_\theta(\gamma) &\equiv \int_{x=0}^\infty \mathrm{e}^{-\gamma\cdot x}\,\mathrm{d}\{\theta(x)\} = \gamma \int_{x=0}^\infty \theta(x)\,\mathrm{e}^{-\gamma\cdot x}\,\mathrm{d}x \\
&= \gamma E \int_{x=0}^\infty D^{-1}(x)\,\mathrm{e}^{-\gamma\cdot x}\,\mathrm{d}x = \gamma\, E \int_{x=0}^\infty \left(\int_{t=0}^\infty 1_{(D(t)\le x)}\,\mathrm{d}t\right)\mathrm{e}^{-\gamma\cdot x}\,\mathrm{d}x \\
&= \int_{t=0}^\infty E\left(\int_{x=0}^\infty 1_{(D(t)\le x)}\,\gamma\,\mathrm{e}^{-\gamma\cdot x}\,\mathrm{d}x\right)\mathrm{d}t = \int_{t=0}^\infty \left(E\,\mathrm{e}^{-\gamma D(t)}\right)\,\mathrm{d}t \\
&= \int_{t=0}^\infty \mathcal{L}_{D(t)}(\gamma)\,\mathrm{d}t = \int_{t=0}^\infty [\mathcal{L}_{D(1)}(\gamma)]^t\,\mathrm{d}t = \frac{-1}{\ln[\mathcal{L}_{D(1)}(\gamma)]}\,.
\end{aligned} \tag{8}$$

The last two equalities follow from (4) and (5).

The gamma-distributed Levy demand process

We now turn our attention to a specific Levy demand process.

Definition 2. The *gamma-distributed Levy process* is the Levy demand process $D^*(t)$ for which $D^*(1)$ has an exponential distribution with mean one.

Note that $D^*(t)$ has a gamma distribution with shape parameter t and rate parameter 1. This is without loss of generality; in continuous-time inventory models we can choose our units of measure for time and inventory so that the demand that occurs in one day has a mean of 1 and a variance of 1. This process is not a new one (see, for example, Prabhu [22] and Feller [11], p. 567). Note that

$$\mathcal{L}_{D^*(t)}(\gamma) = \frac{1}{(1+\gamma)^t}\,. \tag{9}$$

Lemma 4. *For the gamma-distributed Levy demand process $D^*(t)$,*

$$\psi^+(x) = \int_x^\infty \frac{\mathrm{e}^{-y}}{y}\,\mathrm{d}y \quad \text{and} \tag{10}$$

$$\theta'(x) = \int_{t=0}^\infty \frac{x^{t-1}}{\Gamma(t)}\mathrm{e}^{-x}\,\mathrm{d}t, \tag{11}$$

where $\theta'(x)$ is the derivative of $\theta(x)$, and $\theta(0) = 0$.

For the gamma-distributed Levy process, the weighted jump size J has an exponential distribution (see (10) and (3)). The graph of $\psi^+(x)$ in Figure 1 indicates that orders for small quanitites of inventory occur much more often than orders for large quantities, but there is no finite upper bound on the maximum order quantity. By Lemma 3 the same can be said for the overshoot quantities.

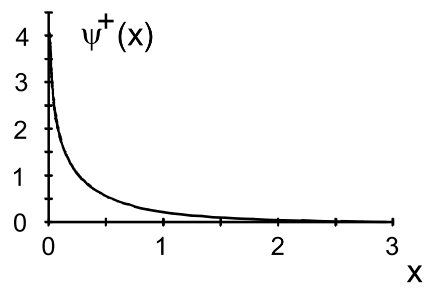

FIGURE 1. $\psi^+(x)$.

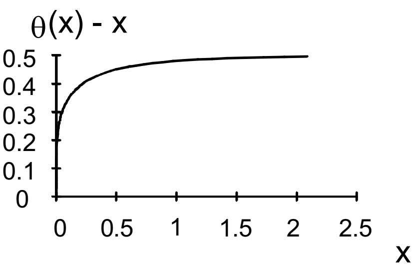

FIGURE 2. $\theta(x) - x$.

A graph of $\theta(x) - x$ appears in Figure 2. From (10) we obtain $E[V] = 1/2$, so by Lemma 2, $\theta(x) - x \to 1/2$ as $x \to \infty$. Suppose that $x > 0.3$. Figure 2 indicates that even though the mean demand rate is equal to one, from a given starting point, the expected amount of time required to accumulate an additional x units of demand is between $x + 0.4$ and $x + 0.5$.

We now list some other properties of the gamma-distributed Levy process $D^*(t)$. The presentation of these properties is designed to facilitate the statements of our lemmas and theorems.

Property 2. $\theta(0^+) = 0$. $\theta(x)$ is non-decreasing and absolutely continuous. $\theta'(x)$ is non-increasing.

Property 3. $\theta(x)$ is continuous, and $\theta'(x) \to 1$ as $x \to \infty$.

Property 4. $\theta(x) > x \cdot \theta'(x)$ for all $x > 0$.

Property 5. $\theta'(x)$ is continuous.

Property 6. For all $x > 0$ there is a $y, 0 < y < x$ such that $\theta'(y) > \theta'(x)$.

Lemma 5. *Properties 1 through 6 hold for $D^*(t)$.*

Note that the first claim of Property 2 implies that $D(t)$ is not a compound Poisson process.

The Mass-Uniform policy, which we develop in Section 4, makes use of the following function.

$$\eta(x) \equiv \int_0^x z \cdot [\theta'(z) - 1] \, \mathrm{d}z. \tag{12}$$

Lemma 6. *For the gamma-distributed Levy process, $\eta(x) \to 1/12$ as $x \to \infty$.*

In optimizing and evaluating (s, S) inventory policies for the gamma-distributed Levy process, the functions $\theta(x)$, $\theta'(x)$ and $\eta(x)$ are frequently used. They were tabulated using numerical integration, and numerically approximated. The approximations are given in Appendix 2. Let IP have the same distribution as the steady-state inventory position. In the next section we will need to compute $E[f(IP)] = \int_0^Q f(S - x)\frac{\theta'(x)}{\theta(Q)}\mathrm{d}x$ numerically for several different functions $f(x)$. The functions $f(x)$ of interest are well-behaved, but the probability measure $\frac{\theta'(x)}{\theta(Q)}\,\mathrm{d}x$ is ill-behaved in the neighborhood of $x = 0$. In fact, for the gamma-distributed Levy process, $\theta'(x) \cdot x\,(\ln(x))^2 \to 1$ as $x \to 0$. Table 1 gives some values of the probability density $\frac{\theta'(x)}{\theta(Q)}$ for $Q = 1$. A transformation described in Appendix 2 was used to perform these integrations in a numerically stable manner.

TABLE 1. The density $\theta'(x)/\theta(1)$.

x	$1\,e-10$	$1\,e-8$	$1\,e-6$	$1\,e-4$	0.01	0.1	1
$\theta'(x)/\theta(1)$	$1.3\,e+7$	$2.1\,e+5$	$3.8\,e+3$	85.9	3.4	1.12	0.70

3. The inventory model

We now turn our attention to (s, S) inventory policies. We consider a continuous-review, single-item inventory system with backorders and with a deterministic lead time L. There is a fixed ordering cost $K > 0$ (in dollars), a holding cost h (in dollars per item per day), and a time-weighted backorder cost p (in dollars per item per day). We choose our units of measure so that the demand $D(1)$ that occurs in one day has a mean and a variance of one. Let D be a random variable whose distribution matches the demand that occurs in one lead time. Thus

$$D \sim D(L) \quad \text{and} \quad E[D] = L.$$

The minimum order quantity is $Q = S - s$. For Levy demand processes orders come in irregular quantities because we usually over-shoot the reorder point. For the gamma-distributed Levy process the actual order quantity is strictly greater than Q with probability one.

Let $NI(t)$ be the net inventory at time t and let $IP(t) \equiv NI(t+L)+D(t, t+L]$ be the inventory position at time t. $IP(t)$ and $NI(t)$ are right-continuous. They have steady-state distributions $F_{IP}(z)$ and $F_{NI}(z)$, where $NI \equiv IP - D$, and D and IP are independent (Zipkin [33]). Since $\theta(x)$ is equal to the mean time required for x units of demand to accumulate after an order is placed, $P(IP \geq S - x) = \theta(x)/\theta(Q)$, $0 \leq x \leq Q$. The average order quantity and the mean time between orders are both equal to $\theta(Q)$. The quantity $\theta(Q) - Q$, referred to in Figure 2, is the mean quantity by which the demand process overshoots the reorder point.

Recall that $\mathrm{E}[D] = L$. If the inventory position at time t is $IP(t) = x$ then the expected rate at which holding and backorder costs will be incurred at time $t + L$ is

$$G(x) \equiv E[h(x - D)^+ + p(D - x)^+] = h(x - L) + (h + p) \cdot n_D(x) \tag{13}$$

where $n_D(x) = E[(D - x)^+]$ is the partial expectation. Note that $G(x)$ is continuous and convex, and that

$$-p \leq G'(x) \leq h, G'(x) \to h \text{ as } x \to \infty, \text{ and } G'(x) \to -p \text{ as } x \to -\infty. \tag{14}$$

Let

$$w \equiv \sup\{x : G'(x) < h\} \cdot$$

If D has a gamma distribution then $w = \infty$ and the following properties hold.

Property 7. $G(x)$ is convex, is strictly convex for $0 < x < w$, and $G'(x) < 0$ for $x < 0$.

Property 8. $G'(x)$ is continuous, and $G''(x)$ is continuous except possibly at $x = 0$.

Following Zheng [32] and Zipkin [33], the expected holding and backorder cost per cycle incurred by an (s,S) policy with $Q = S - s$ is

$$H(S,Q) \equiv \theta(Q) \cdot E[G(IP)] = \int_0^Q G(S-x)\theta'(x)\,\mathrm{d}x. \tag{15}$$

The average cost per day incurred by an (s,S) policy is therefore

$$c(S,Q) = \frac{K + H(S,Q)}{\theta(Q)} \cdot \tag{16}$$

Following the notation in (15), we define

$$H'(S,Q) \equiv \int_0^Q G'(S-x)\ \theta'(x)\,\mathrm{d}x \tag{17}$$

$$H''(S,Q) \equiv \int_0^Q G''(S-x)\ \theta'(x)\,\mathrm{d}x \quad \text{and} \tag{18}$$

$$E(S,Q) \equiv \theta(Q) \cdot G(S-Q) - K - H(S,Q). \tag{19}$$

The first-order optimality conditions are

$$0 = \theta(Q) \cdot \frac{\partial\, c(S,Q)}{\partial\, S} = H'(S,Q) \quad \text{and} \tag{20}$$

$$0 = \frac{\theta(Q)}{\theta'(Q)} \cdot \frac{\partial\, c(S,Q)}{\partial\, Q} = G(S-Q) - c(S,Q) = \frac{E(S,Q)}{\theta(Q)} \cdot \tag{21}$$

Note that (21) states that $G(S-Q)$ is the average cost, which corresponds exactly with Zheng [32]. Zheng's other optimality condition is $G(S-Q) = G(S)$, a simpler, special case of (20).

Lemma 7. *Suppose that S is chosen optimally for a given $Q > 0$. Properties 2 and 3 imply that $G(S-Q) \geq G(S) \geq G(S-Q) - h \cdot [\theta(Q) - Q\,]$. If Properties 4 and 7 also hold then $G(S-Q) > G(S)$.*

Recall that $\theta(Q) - Q \to E[V]$ as $Q \to \infty$. We claim that if S is chosen optimally for a given Q then

$$G(S-Q) - G(S) \ \to\ h \cdot E[V] \ \text{ as } Q \to \infty.$$

If S and Q are large then (20), Properties 2 and 3, equation (14), and Lemma 2 indicate that $G(S-Q) = G(S) + \int_0^Q G'(S-x)[\theta'(x)-1]\,\mathrm{d}x \approx G(S) + \int_0^Q h \cdot [\theta'(x)-1]\,\mathrm{d}x = G(S) + h \cdot [\theta(Q) - Q\,] \approx G(S) + h \cdot E[V]$. A rigorous argument can be constructed along similar lines.

We define y_0 by

$$G'(y_0) = 0. \tag{22}$$

$G(y_0)$ is often referred to as the newsvendor cost or the buffer cost, because when the order cost is equal to zero the optimal policy has $Q = 0$ and $S = y_0$, and $G(y_0)$ is the average cost of that policy.

By the convexity of $G(x)$ and by (20), $S \geq y_0 \geq S - Q$.

Lemma 8. *Suppose that S is chosen optimally for a given $Q > 0$. If Properties 7 and 8 hold then*

$$S > y_0 > S - Q. \tag{23}$$

If Property 2 also holds, the solution to the first-order optimality conditions (20, 21) is unique.

Algebraically the first-order optimality conditions are most conveniently expressed in terms of S and Q. However cost minimization is most efficiently carried out in terms of $s \equiv S - Q$ and S. Let $C(s, S) \equiv c(S, S - s)$. Then

$$\frac{\partial C(s,S)}{\partial S} = \frac{\partial c(S,Q)}{\partial S} + \frac{\partial c(S,Q)}{\partial Q}\bigg|_{Q=S-s} \quad \text{and} \quad \frac{\partial C(s,S)}{\partial s} = -\frac{\partial c(S,Q)}{\partial Q}\bigg|_{Q=S-s}. \tag{24}$$

By (21) and (24),

$$\begin{aligned} \frac{\partial^2 C(s,S)}{\partial s \partial S} &= \left(\frac{\partial}{\partial S} + \frac{\partial}{\partial Q}\right)\left(-\frac{\partial c(S,Q)}{\partial Q}\right)\bigg|_{Q=S-s} \\ &= \frac{\theta'(Q)}{\theta(Q)}\left[\frac{\partial c(S,Q)}{\partial S} + \frac{\partial c(S,Q)}{\partial Q}\right] \\ &\quad + [\, c(S,Q) - G(S-Q)]\left[\frac{\partial}{\partial Q}\left(\frac{\theta'(Q)}{\theta(Q)}\right)\right]\bigg|_{Q=S-s}, \end{aligned} \tag{25}$$

which is equal to zero whenever the first-order conditions are satisfied. This fact suggests that if we alternate Newton steps in s and S, we are likely to get overall quadratic convergence. What we actually do is similar, and it works for the same reason, but the algebra is somewhat simpler. We alternate between Newton steps in S which attempt to find a zero of $H'(S, S - s)$ in (s, S)-space, and Newton steps in s which attempt to find a zero of $E(S, S - s)$ in (s, S)-space. Because the Newton steps are performed in (s, S)-space rather than (S, Q)-space, the derivative of $H'(S, S - s)$ with respect to S and the derivative of $E(S, S - s)$ with respect to s are computed as in (24). The algorithm for systems with discrete demands described by Federgruen and Zheng [7] is similar in that it also alternates between improvements in s and improvements in S.

Cost minimization algorithm

Step 1. Select an initial policy (s, S).

Step 2. Set $s' \leftarrow s - \frac{E(S, S-s)}{G'(s)\,\theta(S-s)}$.

Step 3. Set $S' \leftarrow S - \frac{H'(S,S-s')}{H''(S,S-s')+G'(s')\,\theta'(S-s')}$.

Step 4. Set $s \leftarrow s'$ and $S \leftarrow S'$.

Step 5. Either terminate or go to Step 2.

Lemma 9. *The Cost Minimization Algorithm is quadratically convergent if Properties 2, 7 and 8 hold, and if* $\theta(x)$ *is twice continuously differentiable in a neighborhood of the optimal* Q.

The proofs of Lemmas 7 through 9 assume that D and $\theta(x)$ have certain properties, but they do not assume that they are related to a common Levy demand process in the sense of Section 2. In particular, if we assume that $\theta(x) \equiv x$ and allow D to be arbitrary, Lemmas 7 through 9 hold. Under these assumptions our computation of (s,S) policies corresponds exactly with Zheng's computation of (R,Q) policies. The Cost Minimization Algorithm is therefore quadratically convergent for Zheng's (R,Q) policies, if Properties 7 and 8 hold.

Service levels

We conclude this section by discussing service levels. In discrete time inventory models the Type 1 Service Level is defined as the probability that a time period ends without any backorders. The continuous-time analogue is $P[NI \geq 0]$, the fraction of the time that the net inventory is non-negative. The Type 2 Service Level is the fill rate, the fraction of the demand that is met without backorders. In many continuous-time inventory models, including that of Zheng [32], these two service measures are equivalent. However they are not equivalent for Levy demand processes.

We compute the Type 1 Service Level $P[NI \geq 0]$ as follows. By (13), $G'(x) = h - (h+p)\bar{F}_D(x)$. Hence

$$\begin{aligned} H'(S,Q) &\equiv \int_0^Q G'(S-x)\ \theta'(x)\,\mathrm{d}x \\ &= h \cdot \theta(Q) - (h+p)\int_0^Q \bar{F}_D(S-x)\,\theta'(x)\,\mathrm{d}x \\ &= \theta(Q)\ [h-(h+p)\,P\,[D > IP]] \\ &= \theta(Q)\ [h-(h+p)[1-\ P\,[NI \geq 0]]\,. \end{aligned}$$

We have proven the following lemma, which is reminiscent of Zheng's fill-rate formula.

Lemma 10. *In* (s,S) *policies for Levy demand processes,*

$$P[NI \geq 0] = \frac{p}{p+h} + \frac{H'(S,Q)}{(h+p)\,\theta(Q)},$$

which is equal to $p/(p+h)$ *for an optimal policy.*

Lemma 11. *For Levy demand processes, the fill rate of an* (s,S) *policy is*

$$1 - E\big[(J - (IP - D)^{+})^{+}/J\big]$$

where J *is the weighted jump size (see Lem. 1), and* J, IP, *and* D *are independent.*

Lemma 11 can be explained intuitively as follows. Suppose that a demand of size $D(t) - D(t^{-})$ occurs at time t. At time t^{-} the on-hand inventory is $[IP(t-L) - D(t-L\,,\,t^{-})]^{+}$. The number of units ordered at time t, but not delivered to the client at time t, is

$$\Big\{[D(t) - D(t^{-})] \,-\, \big[IP(t-L) - D(t-L\,,\,t^{-})\big]^{+}\Big\}^{+}\,. \tag{26}$$

Note that the random variables $[D(t) - D(t^{-})]$, $IP(t-L)$ and $D(t-L\,,\,t^{-})$ are stochastically independent. $D(t-L\,,\,t^{-}) \sim D$, $IP(t-L)$ has steady-state distribution $F_{IP}(z)$, and the weighted jump size satisfies $[D(t) - D(t^{-})] \sim J$ (see Lem. 1). Thus the steady-state version of (26) is the formula in Lemma 11.

It is interesting to compare these two service measures, $P(NI > 0)$ and the fill rate. Note that $1_{(NI \le 0)} \le (J - NI^{+})^{+}/J \le 1_{(NI \le J)}$. If β is the fill rate then taking expectations leads to $\mathrm{P}(NI > 0) \ge \beta \ge \mathrm{P}(NI > J)$. If either the lead time L or the order quantity Q is sufficiently large, the standard deviation of $NI = IP - D$ will be large relative to the jump size J, and $P(NI > J) - P(NI > 0)$ will be close to zero.

On the other hand these two service measures can be very far apart. Suppose that the order cost K, the order quantity Q, the lead time L and the order-up-to level S are all very small. Then the jump size J will be large relative to NI, and it is possible for $P(NI < 0)$ to be close to zero and $E[(J - NI^{+})^{+}/J]$ to be close to one. This would correspond to a "veneer inventory" policy, in which only very small orders can be filled from stock. A large fraction of the total demand comes in large orders that must be backordered. The backorder costs are time-weighted and the lead time is assumed to be short relative to the average delay between consecutive large orders. Therefore the average backorder cost incurred is small, and the policy can be economical.

For the gamma-distributed Levy process, the computation of the fill rate requires a two-dimensional numerical integration, assuming that $F_D(x)$ is evaluated using efficient approximations.

4. Policies: Description and analysis

We consider four different policies for this problem, two policies that are based on the classical Hadley–Whitin inventory model, the (s,S) policy whose parameters are taken from Zheng's (R,Q) policy, and a new policy designed for Levy Jump Processes that we call the Mass Uniform Policy.

The Hadley–Whitin inventory model is currently the mainstay of introductory inventory courses. It differs from our inventory model (and from Zheng's) in that it uses a cost of $\hat{p}$ dollars per item for backorders, rather than p dollars per item per day. The change in units of measure complicates direct comparison of the models. We will compare them by following the common practice of using service level targets to determine the backorder costs.

In discussing the Hadley–Whitin model we use $\lambda = 1$ to represent the mean rate at which demand occurs. The Hadley–Whitin cost function for an (s, S) policy with $Q = S - s$ is

$$K\lambda/Q \; + h\,(Q/2 + s - \lambda\,L) + (\hat{p}\,\lambda\,n_D(s))/Q \tag{27}$$

(see, for example, Nahmias [21], p. 258). If the demand process can over-shoot the reorder point, the first term is an approximation. The second and third terms are also approximations. The third term over-states the marginal benefit of an extra unit of safety stock and the second term over-states the marginal cost of an extra unit of safety stock. Although this cost expression never has a global minimum (let $Q > \hat{p}\lambda/h$ and let $s \to -\infty$), it usually has an easily-computed local minimum that corresponds to a very effective policy. However when K (and consequently Q) is sufficiently large, the error in the holding cost term dominates and (27) fails to have a local minimum.

There are two standard computational approaches to this model – the cost minimization approach and the service-constrained approach. The cost-minimization approach attempts to minimize (27) directly. The service-constrained approach searches for a policy that meets a specified service target, and that is optimal for some $\hat{p}$. We lack a backorder cost $\hat{p}$, but we know that for optimal policies, the fraction of time that we have inventory on hand is equal to $p/(p+h)$. According the Hadley–Whitin model this is equivalent to the fill rate, and is equal to $1 - n_D(s)/Q$. So we follow the standard service-constrained approach and obtain a policy such that

$$1 - n_D(s)/Q \; = p/(p + h). \tag{28}$$

Setting the derivative of (27) with respect to s equal to zero, we obtain

$$h - \frac{\hat{p}\,\lambda}{Q}\,\bar{F}_D(s) \; = \; 0. \tag{29}$$

Similarly, the value of Q at a local minimum of (27) must be

$$Q = \sqrt{2(K + \hat{p}\,n_D(s))\,\lambda/h}\,. \tag{30}$$

Substituting (30) into (27) we obtain an average cost of

$$\sqrt{2h(K + \hat{p}\,n_D(s))\,\lambda} \; + h(s - \lambda\,L\,). \tag{31}$$

The derivative of (31) with respect to s converges to $h > 0$ as $|s| \to \infty$. If the derivative ever becomes negative, the supremum of all s for which the derivative is negative is a local minimum of (27). The corresponding policy satisfies (29), and the policy can be obtained through either the cost-minimization approach or the service-constrained approach. However the derivative of (31) may never become negative, in which case the cost minimization approach will fail to produce a policy (in (29), $hQ/\hat{p}\,\lambda > 1$). The service-constrained approach is more robust. It never fails to produce a policy which meets the target service level, and it produces policies which cannot be obtained using the cost-minimization approach.

We study two policies that are based on the Hadley–Whitin model. Our first policy, called HW-COST, is the standard service-constrained algorithm (see, for example, Nahmias [21], pp. 263-264). We eliminate $\hat{p}$ from (29) and (30), and the resulting equation is solved together with (28). This algorithm fails only when $h \geq p$. The standard cost minimization algorithm, which is equivalent when it works, but which works less often, is as follows. Guess at $\hat{p}$ and find the local minimum of (27). Then search for a $\hat{p}$ such that the computed policy satisfies $1 - n_D(s)/Q = p/(p+h)$. It may be that $1 - n_D(s)/Q > p/(p+h)$ for all values of $\hat{p}$ for which (27) has a local minimum. In that case we say that HW-COST has failed, meaning that to obtain the desired service measure we would need to make $\hat{p}$ small enough that (27) would not have a local minimum. We compute and evaluate the policy whether this happens or not.

Our second policy, called HW-EOQ, is the common approach of using the EOQ model to select Q and using the service target to select s. Thus $Q = \sqrt{2\,K\,\lambda/h}$ and s is selected so that (28) holds (see, for example, Nahmias [21], p. 262). If $s \geq 0$ then (29) gives an imputed $\hat{p}$. If $s < 0$ we say that HW-EOQ has failed. In this case, if $Q = \sqrt{2\,K\,\lambda/h}$ then the standard cost-minimization algorithm will fail ($hQ/\hat{p}\,\lambda > 1$ in (29)), and all values of $\hat{p}$ for which (28) has a local minimum yield policies with service levels that are higher than the target.

Our third policy, called ZHENG, is the (s, S) policy whose parameters are taken from Zheng's (R, Q) policy (Zheng [32]). The policy is computed using the Cost Minimization Algorithm with $\theta(x) \equiv x$.

Our fourth policy is called the Mass Uniform policy, or MASS-U. For a given $\overline{Q} > 0$ we approximate $\theta(x)$ with a function $\mu(x)$ defined as follows.

$$\begin{aligned}
q(\overline{Q}) \;&\equiv \int_0^{\overline{Q}} \left(\theta'(x) - \theta'(\overline{Q})\right)\, \mathrm{d}x \;=\; \theta(\overline{Q}) - \overline{Q}\cdot\theta'(\overline{Q}),\\
a(\overline{Q}) \;&\equiv \frac{1}{q(\overline{Q})} \int_0^{\overline{Q}} x\cdot[\theta'(x) - \theta'(\overline{Q})]\, \mathrm{d}x \;=\; \left[\eta(\overline{Q}) - \frac{1}{2}\,\overline{Q}^{\,2}\left(\theta'(\overline{Q}) - 1\right)\right]/q(\overline{Q}),\\
\mu(x) \;&\equiv\; \theta'(\overline{Q})\cdot x \,+\, q(\overline{Q})\cdot 1_{(x \geq a(\overline{Q}))}\,.
\end{aligned} \tag{32}$$

The dependence of $\mu(x)$ on $\overline{Q}$ is suppressed. The measure $\mu(x)$ takes the area that lies under the curve $\theta'(x)$, $0 \leq x \leq \overline{Q}$ and above $\theta'(\overline{Q})$, and concentrates it into a mass of size $q(\overline{Q})$ located at $x = a(\overline{Q})$ (see Fig. 3). $a(\overline{Q})$ is chosen to equalize the first moments of the measures $\theta'(x)\,\mathrm{d}x$, $0 \leq x \leq \overline{Q}$ and $\mathrm{d}\{\mu(x)\}$, $0 \leq x \leq \overline{Q}$.

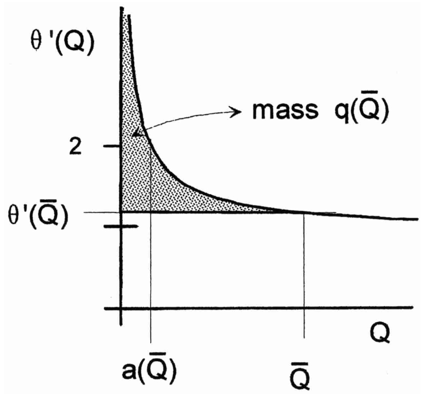

FIGURE 3. The mass uniform policy.

Since $[\theta'(x) - \theta'(\overline{Q})]/q(\overline{Q}),\ 0 \le x \le \overline{Q}$ is a probability density, equation (32) and Properties 2 and 6 imply that

$$0 < q(\overline{Q}) \text{ and } 0 < a(\overline{Q}) \le \overline{Q}/2 \quad \text{for all } \ \overline{Q} > 0. \tag{33}$$

If we use $\mathrm{d}\{\mu(x)\}$ in place of $\theta'(x)\,\mathrm{d}x$, and if $Q \ge a(\overline{Q})$, then (16) becomes

$$c^{\overline{Q}}(S,Q) \equiv \left[K + q(\overline{Q}) \cdot G(S - a(\overline{Q})) + \theta'(\overline{Q}) \int_0^Q G(S-x)\ \mathrm{d}x\right] / \left[q(\overline{Q}) + Q \cdot \theta'(\overline{Q})\right]. \tag{34}$$

Consider the optimization problem

$$(P^{\overline{Q}}) \qquad \begin{aligned} &\text{min: } c^{\overline{Q}}(S,Q) \\ &\text{such that: } Q \ge a(\overline{Q}). \end{aligned}$$

Let $S(\overline{Q})$ and $Q(\overline{Q})$ solve (P$^{\overline{Q}}$), and suppose that $a(\overline{Q}) < Q(\overline{Q})$. Lemma 9 implies that the Cost Minimization algorithm converges quadratically to $(S(\overline{Q}), Q(\overline{Q}))$. The Cost Minimization Algorithm is easily adapted to solve (P$^{\overline{Q}}$).

To compute the Mass Uniform policy we select a $\overline{Q}$ and we use the Cost Minimization Algorithm with appropriate modifications to solve $(P^{\overline{Q}})$. We iteratively adjust $\overline{Q}$ to make $\overline{Q} = Q\,(\overline{Q})$. The approximations in Appendix 2 are helpful. Unlike the computation of optimal policies, equation (32) allows us to avoid numerical integrations. The following lemma guarantees that the approach works.

If we use $\mathrm{d}\{\mu(x)\}$ in place of $\theta'(x)\,\mathrm{d}x$, and if $Q \geq a(\overline{Q})$, then (16) becomes

$$c^{\overline{Q}}(S,Q) \equiv \left[K + q(\overline{Q}) \cdot G(S - a(\overline{Q})) + \theta'(\overline{Q}) \int_0^Q G(S-x)\ \mathrm{d}x\right] / \left[q(\overline{Q}) + Q \cdot \theta'(\overline{Q})\right]. \tag{35}$$

Consider the optimization problem

$$(P^{\overline{Q}}) \qquad \min\colon\ c^{\overline{Q}}(S,Q) \qquad \text{such that: } Q \geq a(\overline{Q}).$$

Let $S(\overline{Q})$ and $Q(\overline{Q})$ solve $(P^{\overline{Q}})$, and suppose that $a(\overline{Q}) < Q(\overline{Q})$. Lemma 9 implies that the Cost Minimization algorithm converges quadratically to $(S(\overline{Q}), Q(\overline{Q}))$. The Cost Minimization Algorithm is easily adapted to solve $(P^{\overline{Q}})$.

To compute the Mass Uniform policy we select a $\overline{Q}$ and we use the Cost Minimization Algorithm with appropriate modifications to solve $(P^{\overline{Q}})$. We iteratively adjust $\overline{Q}$ to make $\overline{Q} = Q(\overline{Q})$. The approximations in Appendix 2 are helpful. Unlike the computation of optimal policies, equation (32) allows us to avoid numerical integrations. The following lemma guarantees that the approach works.

Lemma 12. *Assume that Properties 2, 3, 5, 6, 7 and 8 hold and that* $\overline{Q} > 0$. *The optimal solution* $(S(\overline{Q}), Q(\overline{Q}))$ *of* $(P^{\overline{Q}})$ *is the unique solution of the first-order optimality conditions for* $(P^{\overline{Q}})$. $S(\overline{Q})$ *and* $Q(\overline{Q})$ *are continuous functions of* $\overline{Q}$. *If* $\overline{Q}$ *is sufficiently large then* $\overline{Q} > Q(\overline{Q})$, *and if* $\overline{Q} > 0$ *is sufficiently small then* $\overline{Q} < Q(\overline{Q})$.

The relative cost of a policy is defined to be $(c' - c^*)/c^*$ where c' is the average cost of the policy and c^* is the average cost of an optimal policy. Optimal policies are computed using use the Cost Minimization Algorithm.

Lemma 13. *The relative cost of Zheng's policy is at most* $\min\left\{\frac{\theta(Q)-Q}{Q+G(y_0)/p}, \frac{h+p}{p}\frac{[\theta(Q)-Q]^2+\eta(Q)}{Q^2}\right\}$, *where* $Q = S - s$ *is the optimal order quantity (not Zheng's order quantity).*

For the gamma-distributed Levy process, $\theta(Q) - Q \to \mathrm{E}[V] = 1/2$ and $\eta(Q) \to 1/12$ as $Q \to \infty$, so the relative cost of Zheng's policy is at most $\min\left\{\frac{0.5}{Q+G(y_0)/p}, \frac{h+p}{p}\frac{1}{3\,Q^2}\right\}$.

Note that $G(y_0) \to 0$ as $L \to 0$, and $G(y_0) \to \infty$ as $L \to \infty$. The bound of Lemma 13 implies that if $L > \epsilon > 0$ then Zheng's policy works very well as $Q \to 0$, as $Q \to \infty$ and as $L \to \infty$. However there is no uniform, finite upper bound on the relative cost of Zheng's policy (proof omitted), and the trends illustrated by this bound are similar to the computational results of Section 5. If the lead time is zero or close to zero, and the order cost K is small, Zheng's policy can be far from optimal. This statement probably applies to all existing

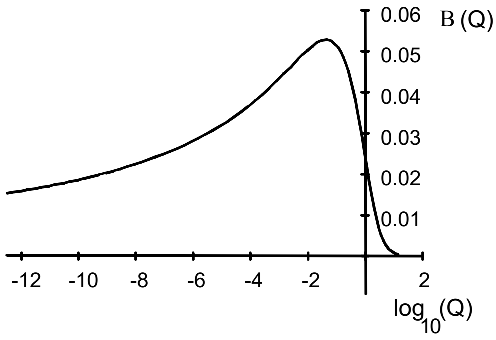

FIGURE 4. The relative cost of the mass uniform policy.

continuous-time (s, S) inventory models which assume the inventory position to be uniformly distributed.

Lemma 14. *The relative cost of the Mass Uniform Policy* (S, Q) *is at most* $\frac{h+p}{p} B(Q)$*, where*

$$B(Q) \equiv \frac{1}{Q \cdot \theta(Q)} \int_0^Q x \cdot (\theta'(x) - \theta'(Q))\, \mathrm{d}x \;=\; \frac{q(Q) \cdot a(Q)}{Q \cdot \theta(Q)} \,.$$

For the gamma-distributed Levy process, the function $B(Q)$ is graphed in Figure 4. Since $B(Q)$ achieves its maximum value of 0.0527 at $Q = 0.0307$, the relative cost of the Mass Uniform policy is at most 0.0791 if $p/h \geq 2$. Since $B(1) = 0.0235$, the relative cost of the Mass Uniform policy is at most 0.0353 if $p/h \geq 2$ and $Q \geq 1$.

5. Computational results

The main purpose of our computational study is to use the gamma-distributed Levy process as a vehicle for testing the robustness of the (s, S) inventory policies described in Section 4. The bound in Lemma 14 is strong enough that MASS-U does not need computational validation, but the performance that can be expected of the other policies is less certain. In addition we want to gain intuition into the nature of optimal policies for the gamma-distributed Levy process and to explore the importance of modelling the overshoot and the non-uniformity of the distribution of the inventory position.

As has been mentioned before, we scale our units of measure for time and inventory so that the demand which occurs in one time period has a mean and a variance of one. In this section we select our unit of measure for money so that the holding cost is $h = 1$. The fact that we re-scaled out units of measure alters the intuitive meaning of the remaining parameters. The backorder cost p is interpreted as a measure of service. Lemma 10 implies that $1/(1 + p)$ is the fraction of time that we are out of stock. L can be thought of as a measure of the variability of the demand that occurs in one lead time. In particular, D has a squared coefficient of variation of $1/L$. The order cost K is a prime determinant of the minimum order quantity Q. We define the *reorder interval* to be Q, the time interval corresponding to the minimum order quantity. The squared coefficient of variation of the demand that occurs in a reorder interval is $1/Q$.

Application range

As a vehicle for interpretating our results we define the "application range", a domain of the parameter space in which most real-world applications of (s, S) inventory systems lie. Since this is primarily based on personal experience it is bound to be somewhat controversial. In our experience lead times can be long or short, so all lead times are included in the application range. Most inventory systems operate with at least moderately high service levels, so for membership in the application range we require $p \geq 3$ (the system is out of stock at most 25% of the time; see Lem. 10).

We define the order costs that lie in the application range indirectly, through the order quantities Q. In our experience most large and moderately large inventory systems have a substantial number of items which experience spikey or sporadic demand. For these items the squared coefficient of variation of the demand that occurs in Q days is often very high. Maintaining inventories for these products tends to be very expensive. In a great many cases distribution systems should be re-designed to eliminate the need to inventory these parts, but this is not always possible. Our application range is intended to include some parts of this type, but not all of them.

We considered the following criteria for an order quantity that falls within the application range. First, the mean order quantity is at most twice the minimum order quantity, *i.e.*, $\theta(Q^Z) \leq 2Q^Z$, where Q^Z comes from ZHENG. (This translates into $Q^Z \geq 0.444$. Taking Q from HW-COST or HW-EOQ gives results that are comparable, but somewhat different.) Second, the demand that occurs in Q^Z days has a squared coefficient of variation of at most 2. (This translates into $Q^Z \geq 0.5$.) Third, the demand that occurs in $\theta(Q^O)$ days has a squared coefficient of variation of at most 2, where Q^O comes from the optimal policy. (This translates into $Q^O \geq 0.134$. Taking Q from MASS-U gives nearly identical results.)

In the context of our computational studies these three criteria turned out to be nearly equivalent. We chose the first one. Thus a problem instance is said to fall within the application range if $p \geq 3$ and if $Q^Z \geq 0.444$, where Q^Z comes from ZHENG.

Results

Figure 5 illustrates how the relative order quantities change as the order cost decreases. For $K \geq 0.25$ it appears that the Q values for ZHENG approximate the optimal value of $\theta(Q)$. This trend has intuitive appeal because the average order quantity is $\theta(Q)$ for these processes, and it is Q in Zheng's paper. However the trend breaks down for smaller order quantities. Qualitatively, HW-COST and HW-EOQ behave similarly to ZHENG, MASS-U behaves like the optimal policy, and the backorder cost and the lead time have little or no impact.

Figure 6 illustrates the fact that the relative cost of ZHENG matches the qualitative behavior of the theoretical bound in Lemma 13. ZHENG, HW-COST and HW-EOQ all have have unbounded relative cost for $L = 0$ and $K \approx 0$. This is true of all policies for which, if we set $L = 0$, $Q\,/\sqrt{2K\lambda/h}$ fails to converge to 0 as $K \to 0$, including both HW-COST and HW-EOQ. For a fixed $L > 0$ the relative cost of these three policies (and most other (s,S) policies in the literature) converges to 1 as $K \to 0$.

These observations can be explained as follows. Because of the non-uniformity of the distribution of the inventory position, in optimal policies the order-up-to level S is closer to y_0 than it otherwise would be. As K gets small this trend becomes more pronounced. In addition, the expected overshoot grows relative to the minimum order quantity Q, effectively reducing the average order cost incurred per day. Figure 5 confirms this by showing that for ZHENG, Q is too large when the order costs are small. If the lead time is zero or very small these errors can be very costly. On the other hand, if the lead time is positive then the costs of all of these policies converge to the Newsvendor Cost $G(y_0)$ as $K \to 0$, so the relative cost tends to 0.

Our main computational experiment used all combinations of order costs $K =$ 0.0625, 0.25, 1, 4, 16, 64, 256, 1024, 4096, 16384, backorder costs of $p = 1$, 2, 4, 8, 16, 32, 64, 128, 256, 512, and lead times $L = 0$, 0.0625, 0.25, 0.5625, 1, 1.5625, 2.25, 3.0625, 4, 5.0625. Since the HW-COST algorithm does not work for $p \le h = 1$, we substituted $p = 1.2$ for $p = 1$ for the HW-COST policy. In all 1 000 parameter sets were tested. Of the 1 000 parameter sets, 792 are inside of the application range. 200 parameter sets have backorder costs p that are less than 3. For 9 parameter sets ZHENG produces Q values that are less than 0.444, and both of these criteria apply to one parameter set. Note that Figure 6 includes data sets with smaller order costs than our main experiment, but only for $p = 10$.

Table 2 summarizes the results. MASS-U consistently out-performs the theoretical bound given by Lemma 14, usually by a very substantial margin. It was never more than 3.2% from optimal, and its average relative cost was negligible. On average ZHENG was only 1.3% from optimal, but it was off by as much as 52%, and within the application range it was off by as much as 20.6%. Both HW-COST and HW-EOQ perform well on average, but within the application range they both had maximum relative costs of over 27%. If we had defined the applicaiton range *via* $Q^Z \ge 1$ rather than $Q^Z \ge 0.444$ the maximum relative costs would have been much smaller.

Our main computational experiement contains 1 000 parameter sets, but the gaps between parameter values are still large enough to make the maximum errors reported in Table 2 unreliable. For example, in the experiment that generated the data for Figure 6 we included the parameter set $K = 0.125$, $L = 0$, $p = 10$. This parameter set falls within the application range ($Q = 0.52$), and it has relative

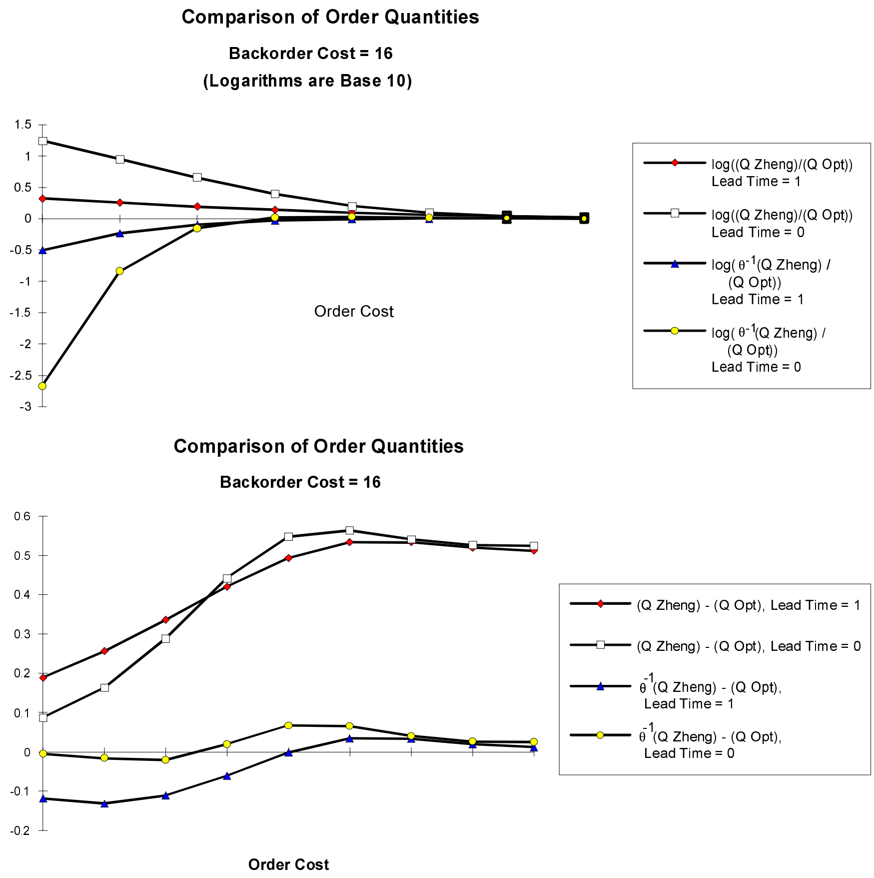

FIGURE 5. Comparison of order quantities.

costs of 40% for HW-Cost, 40% for HW-EOQ, and 33% for Zheng. These numbers are substantially larger than the maximum errors reported in Table 2.

Even within the application range the cost-minimization versions of HW-COST and HW-EOQ fail for over 33% of the parameter sets, because the cost function does not have a local minimum.

Tables 3, 4 and 5 illustrate the combinations of parameters that cause problems for the different policies. Because MASS-U is uniformly very effective no tables were produced for it. The cost-minimization versions of HW-COST and HW-EOQ fail often, both in and out of the application range, especially for larger order costs K, lower backorder costs p, and lower lead times. The HW-COST policy is more than 20% from optimal only when the backorder cost is outside of the application range ($p < 3$), or when both the order cost and the lead time are small. Relatively speaking, HW-EOQ has more trouble when both the order cost and the lead time are small, and is more robust with small backorder costs.

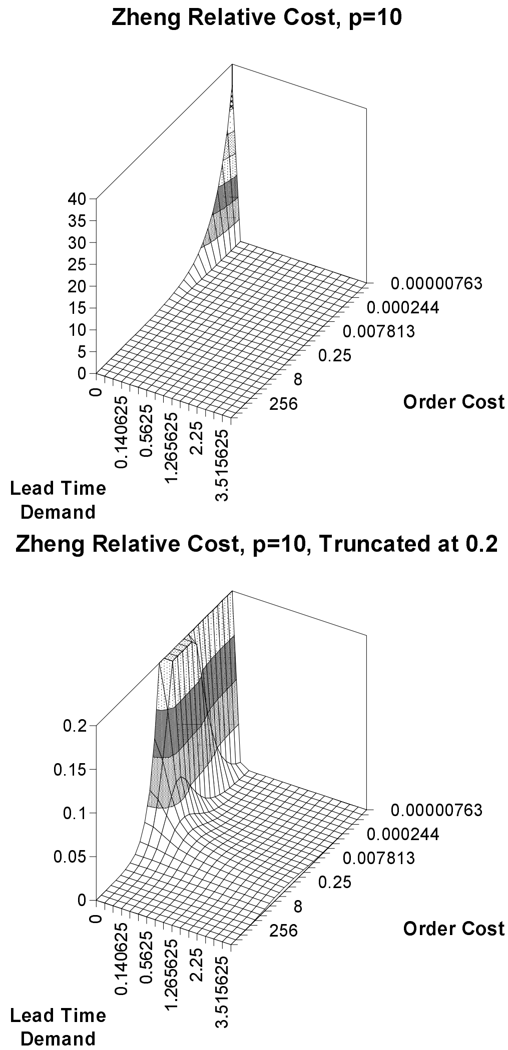

FIGURE 6. Relative costs of the Zheng policy.

ZHENG is more robust than either of the others, but both inside of the application range and outside of it, Zheng's policy has trouble when both the order cost and the lead time are small.

The errors that occur when both the order cost and the lead time are small were explained when we discussed Figure 6. Both the order quantity and the relative cost of the HW-COST policy grow without bound as p approaches $h = 1$. With $p = 1.2$, HW-COST clearly had problems.

Table 2. Relative costs of the policies.

	Relative Cost = (Policy Cost) / (Opt. Cost) −1			
	HW-Cost	HW-EOQ	ZHENG	MASS-U
All Test Problems				
Mean	0.077	0.051	0.013	9.2E-05
St. Dev.	0.166	0.111	0.057	0.001
Max	1.820	0.929	0.523	0.032
% Heur. Fails	0.452	0.411		
In the Application Range (79% of total)				
Mean	0.017	0.030	0.006	4.9E-06
St. Dev.	0.031	0.061	0.023	4.5E-05
Max	0.275	0.392	0.206	0.001
% Heur. Fails	0.343	0.337		
When $Q \geq 1$ and $p \geq 3$				
Max	0.128	0.096	0.070	0.000
When the Heuristic does Not Fail				
Mean	0.025	0.066		
St. Dev.	0.045	0.127		
Max	0.404	0.929		

As Figure 7 illustrates, the ratio of the relative cost of the ZHENG policy to the bound in Lemma 13 is usually much less than one. But in our tests, when the order cost is 16 or more the arithmetic difference between the bound and the relative cost of the policy is at most 1%. Lemma 14 gives a bound on the relative cost of the MASS-U policy. The ratio of this bound to the relative cost of the MASS-U policy is less than 6 in only 2 of the 1 000 parameter sets in the test. Usually it is much higher. However the arithmetic difference between the bound and the relative cost of the policy is never more than 5.2%, and is usually much smaller (see Fig. 7).

Table 6 gives the CPU time per problem instance, in seconds. All policies requiring an initial guess were started from the HW-EOQ policy. For Zheng's policy we used the Cost Minimization Algorithm, which is faster than the algorithm that Zheng proposed. All of the policies have very low computation times except for the optimal policy. The optimal policy requires three numerical integrations for each iteration of the Cost Minimization Algorithm.

Table 3. Performance of the HW-cost policy.

The Largest Lead Time for which HW-Cost Fails											
Lead Times Tested: 0, 0.06, 0.25, 0.56, 1, 1.56, 2.25, 3.06, 4, 5.06											
		Order Costs									
		0.06	0.25	1	4	16	64	256	1024	4096	16384
	1.2	0.25	1	All	All	All	All	All	All	All	All
Back-	2	0.06	0.25	1	1.6	All	All	All	All	All	All
order	4	0.06	0.06	0.25	0.56	1	3.1	All	All	All	All
Costs	8	0	0.06	0.06	0.25	0.56	1	3.1	All	All	All
	16	0	0	0.06	0.06	0.25	0.56	1	3.1	All	All
	32	0	0	0	0.06	0.06	0.25	0.56	1	2.3	All
	64	0	0	0	0	0.06	0.06	0.25	0.56	1	2.3
	128	0	0	0	0	0	0.06	0.06	0.25	0.56	1
	256	0	0	0	0	0	0	0.06	0.06	0.25	0.56
	512	0	0	0	0	0	0	0	0.06	0.06	0.25

Note: When the lead time L is equal to zero, the standard cost-minimization approach leads to a reorder point of 0 and a fill rate of 100%.
Note: In each each row, in the limit as $K \to \infty$, we eventually get ALL.

Lead Times for which HW-Cost has Relative Cost > 20%											
Lead Times Tested: 0, 0.06, 0.25, 0.56, 1, 1.56, 2.25, 3.06, 4, 5.06											
		Order Costs									
		0.06	0.25	1	4	16	64	256	1024	4096	16384
	1.2	All	All	All	All	All	All	All	All	All	All
Back-	2	*0.25	*0.06	None	None	None	None	None	None	None	None
order	4	0	0	None	None	None	None	None	None	None	None
Costs	8	0	0	None	None	None	None	None	None	None	None
	16	0	0	None	None	None	None	None	None	None	None
	32	0	0	None	None	None	None	None	None	None	None
	64	0	0	None	None	None	None	None	None	None	None
	128	0	0	None	None	None	None	None	None	None	None
	256	0	0	None	None	None	None	None	None	None	None
	512	0	0	None	None	None	None	None	None	None	None
Order	Mean	2.2	2.4	2.9	4.2	6.9	12.6	24.0	46.6	92.0	183
Quan-	Min	0.4	0.7	1.4	2.8	5.7	11.3	22.6	45.3	90.5	181
tities	Max	10.1	10.2	10.3	10.9	12.7	17.5	28.2	50.6	95.7	186

TABLE 4. Performance of the HW-cost policy.

The Largest Lead Time for which HW-EOQ Fails
Lead Times Tested: 0, 0.06, 0.25, 0.56, 1, 1.56, 2.25, 3.06, 4, 5.06

		Order Costs									
		0.06	0.25	1	4	16	64	256	1024	4096	16384
	1	0.06	0.25	0.56	1	2.3	All	All	All	All	All
Back-	2	0.06	0.06	0.25	0.56	1.6	3.1	All	All	All	All
order	4	0.06	0.06	0.25	0.56	1	2.3	4	All	All	All
Costs	8	0	0.06	0.06	0.25	0.56	1	2.3	4	All	All
	16	0	0	0.06	0.06	0.25	0.56	1	2.3	All	All
	32	0	0	0	0.06	0.06	0.25	0.56	1	2.3	All
	64	0	0	0	0	0.06	0.06	0.25	0.56	1	2.3
	128	0	0	0	0	0	0.06	0.06	0.25	0.56	1
	256	0	0	0	0	0	0	0.06	0.06	0.25	0.56
	512	0	0	0	0	0	0	0	0.06	0.06	0.25

Note: When the lead time L is equal to zero, the standard cost-minimization approach leads to a reorder point of 0 and a fill rate of 100%.

Lead Times for which HW-EOQ has Relative Cost > 20%
Lead Times Tested: 0, 0.06, 0.25, 0.56, 1, 1.56, 2.25, 3.06, 4, 5.06

		Order Costs									
		0.06	0.25	1	4	16	64	256	1024	4096	16384
	1	All	≥ 1	≥ 4	None	None	None	None	None	None	None
Back-	2	$\neq 0.06$	≥ 1.56	None	None	None	None	None	None	None	None
order	4	$\neq 0.06$	≥ 5.06	None	None	None	None	None	None	None	None
Costs	8	$\neq 0.06$	None	None	None	None	None	None	None	None	None
	16	≥ 0.25, ≤ 2.25	None	None	None	None	None	None	None	None	None
	32	0	None	None	None	None	None	None	None	None	None
	64	0	0	None	None	None	None	None	None	None	None
	128	0	0	None	None	None	None	None	None	None	None
	256	0	0	None	None	None	None	None	None	None	None
	512	0	0	None	None	None	None	None	None	None	None
Order	Mean	0.4	0.7	1.4	2.8	5.7	11.3	22.6	45.3	90.5	181
Quan-	Min	0.4	0.7	1.4	2.8	5.7	11.3	22.6	45.3	90.5	181
tities	Max	0.4	0.7	1.4	2.8	5.7	11.3	22.6	45.3	90.5	181

Table 5. Performance of the Zheng policy.

Lead Times for which Zheng has Relative Cost > 20%

Lead Times Tested: 0, 0.06, 0.25, 0.56, 1, 1.56, 2.25, 3.06, 4, 5.06

		Order Costs									
		0.06	0.25	1	4	16	64	256	1024	4096	16384
	1	≤ 0	≤ 0	None	None	None	None	None	None	None	None
Back-	2	≤ 0	≤ 0	None	None	None	None	None	None	None	None
order	4	≤ 0.06	None	None	None	None	None	None	None	None	None
Costs	8	≤ 0.06	None	None	None	None	None	None	None	None	None
	16	≤ 0.06	None	None	None	None	None	None	None	None	None
	32	≤ 0.06	None	None	None	None	None	None	None	None	None
	64	≤ 0.06	None	None	None	None	None	None	None	None	None
	128	≤ 0.06	None	None	None	None	None	None	None	None	None
	256	≤ 0.06	None	None	None	None	None	None	None	None	None
	512	≤ 0.06	None	None	None	None	None	None	None	None	None
Order	Mean	0.9	1.4	2.3	3.9	7.0	13.0	25.2	49.7	98.8	197
Quan-	Min	0.4	0.7	1.4	2.8	5.7	11.3	22.7	45.3	90.7	181
tities	Max	1.2	1.9	3.2	5.3	9.0	16.6	32.3	64.2	128	256

Range of Order Quantities for the Optimal Policy

Lead Times Tested: 0, 0.06, 0.25, 0.56, 1, 1.56, 2.25, 3.06, 4, 5.06

		Order Costs									
		0.06	0.25	1	4	16	64	256	1024	4096	16384
Order	Mean	0.5	1.0	1.8	3.4	6.4	12.5	24.7	49.2	98.3	197
Quan-	Min	0.1	0.3	0.9	2.3	5.1	10.8	22.1	44.8	90.1	181
tities	Max	0.9	1.6	2.8	4.8	8.5	16.1	31.8	63.7	128	256

Table 6

Computation Times: Avg. Seconds per Problem Instance				
HW-COST	HW-EOQ	ZHENG	MASS-U	Optimal
.027	.007	.076	.069	36.5
Computations done on a 486-based, 50 megahertz PC				

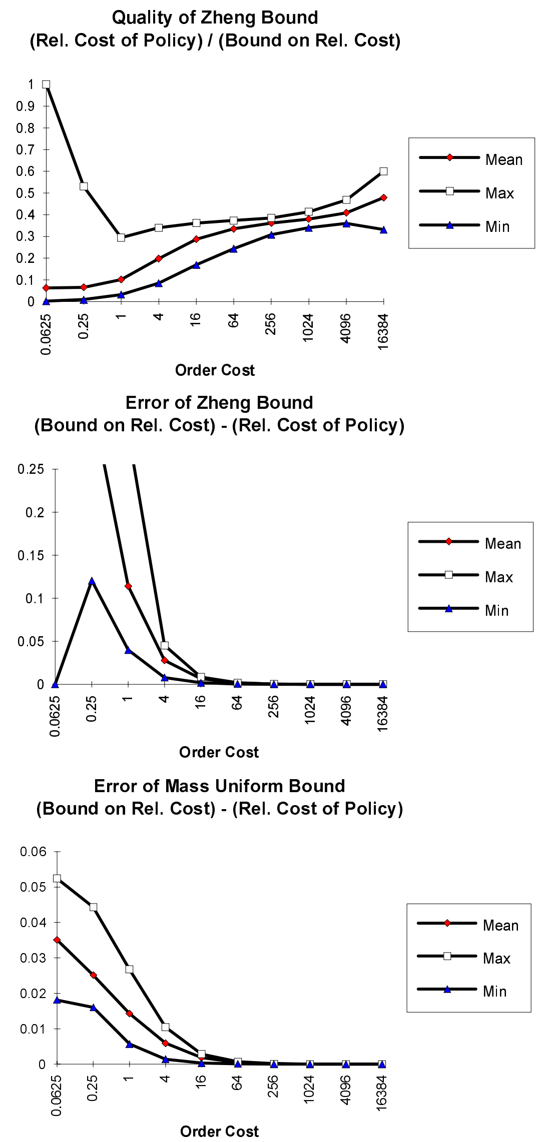

FIGURE 7. Quality of bounds on the relative cost.

RECOMMENDATIONS

Heuristic policies are usually measured on the quality of the policies that they produce and on the computational effort that they require. All of the four heuristic policies studied can be computed very efficiently.

The standard Hadley–Whitin cost function often fails to have a local minimum. However the standard algorithms for computing polices with a given fill rate, Zheng's policy, and the Mass Uniform policy are all much more robust. When the service level is reasonably high ($p/(p+h) \geq 2/3$), and $Q \geq 1$, all of these policies perform very well.

The main negative result of these tests is that when $Q < 1$, as it often is with low-demand parts that experience sproadic or spikey demand, only the MASS-U policy can be relied on. (Recall that because of the way we scaled time and inventory, Q should be interpreted as the coefficient of variation of the demand that occurs in one reorder interval.) Other policies that assume the inventory position to be uniformly distributed would almost certainly experience similar problems. If inventory levels are discrete the algorithm of Federgruen and Zheng [7] should be used. To our knowledge this is the only paper that efficiently computes good (s, S) inventory policies for systems with continuous inventory levels, that does not assume the inventory position to be uniform, and that allows the demand process to over-shoot the reorder point.

6. Conclusions

Levy demand processes are a useful and interesting set of demand processes for inventory models. Numerical approximations for the distribution of D, of θ, and/or both will be required. The gamma-distributed Levy process is particularly attractive, and we have provided the appropriate approximations.

For Levy demand processes the distribution of the inventory position does not need to be uniform, and the demand process is allowed to over-shoot the reorder point. Most easily-computed (s, S) inventory policies require the inventory position to be uniform and assume that there is no overshoot. Our tests indicate that when the coefficient of variation of the demand that occurs in the reorder interval is greater than one, it is important to model the inventory position as non-uniform and to model the overshoot when it occurs. This is often the case for low-demand parts that experience sporadic or spiky demand.

As long as the coefficient of variation of the demand that occurs in one reorder interval is at least one, and the service level is reasonably high, the standard service-constrained Hadley–Whitin (s, S) inventory polices and Zheng's policy work very well. However even in this region it is often the case that the standard Hadley–Whitin cost function fails to have a local minimum.

The Mass Uniform heuristic applies to all Levy demand processes. For the gamma-distributed Levy process it is guaranteed to be within 8% of optimal whenever backorders are at least as expensive as inventory.

For any Levy demand process, the Cost Minimization Algorithm applies to Zheng's (R, Q) inventory model, to the Mass Uniform heuristic and to the computation of optimal policies. The algorithm is quadratically convergent.

Appendix 1: Notation

$a(Q)$	The location of the point mass associated with the Mass Uniform policy (see (32)).
$B(Q)$	$\frac{q(Q)\cdot a(Q)}{Q\cdot\theta(Q)}$ (see Lem. 14).
$c(S,Q)$	The average cost of the (s,S) policy with $S-s=Q$ (see (16)).
$C(s,S)$	$c(S,S-s)$, the average cost of an (s,S) policy.
$D(t,u)$	The demand that occurs in the time interval $(t,u]$.
$D(t)$	The demand that occurs in the time interval $[0,t]$.
D	$D(L)$, the demand that occurs in one lead time.
$E(S,Q)$	$\theta(Q)\cdot G(S-Q)-K-H(S,Q)$ (see (19)).
$\eta(Q)$	$\int_0^x z\cdot[\theta'(z)-1]\ \mathrm{d}z$ (see (12)).
$F_D(x)$	$P(D\le x)$, the cumulative distribution function of the random variable D.
$\bar{F}_D(x)$	$P(D>x)$, the complementary cumulative distribution function of the random variable D.
$G(x)$	$E[h(x-D)^+ + p(D-x)^+]$, the expected rate at which holding and backorder costs are incurred (see (13)).
$G(y_0)$	The Newsvendor Cost or the Buffer cost. y_0 minimizes $G(x)$ (see (22)).
h	The holding cost, in dollars per item per day.
$H(S,Q)$	$\int_0^Q G(S-x)\theta'(x)\ \mathrm{d}x$, the expected holding cost incurred per cycle (see (15)).
$IP(t)$	$NI(t+L)+D(t,\ t+L)$, the inventory position at time t.
J	The Weighted Distribution of the Jump Size (see (3) and Lem. 1).
K	The fixed order cost, in dollars.
L	The lead time.
$\mathrm{L}_D(\gamma)$	$E\left[\mathrm{e}^{-\gamma D}\right]$, the Laplace Transform of the random variable D.
$\mathrm{L}_\theta(\gamma)$	$\int_{x=0}^{\infty}\mathrm{e}^{-\gamma\cdot x}\ \mathrm{d}\{\theta(x)\}$, the Laplace Transform of the function θ.
λ	The demand rate (assumed to be equal to 1).
$n_D(x)$	$E\left[(D-x)^+\right]$, the partial expectation of the random variable D at x.
$NI(t)$	The net inventory at time t.
p	The backorder cost, in dollars per item per day.
Q	The minimum order quantity. $Q=S-s$.
$q(Q)$	The point mass associated with the Mass Uniform policy (see (32)).
s	The reorder point.
S	The order-up-to level.
$\theta(Q)$	$E\left[D^{-1}(Q)\right]$, the expected time to accumulate Q units of demand.
$\theta'(Q)$	The derivative of $\theta(Q)$ (see (11)).
t	Used to index time.
V	The Asymptotic distribution of the overshoot (see (3) and Lem. 3).
y_0	The value that minimizes $G(x)$ (see (22)).

Appendix 2: Numerical approximations of $\theta(x)$ and $\theta'(x)$

In this appendix we describe our approach to the numerical integrations that are needed to compute an optimal policy. We also give the polynomial approximations for $\theta'(x)$, $\theta(x)$ and $\eta(x)$ for the gamma-distributed Levy process, which we developed and used.

Numerical integrations

We need to evaluate the integrals $\int_0^Q G(x)\,\theta'(x)\,\mathrm{d}x$, $\int_0^Q G'(x)\,\theta'(x)\,\mathrm{d}x$ and $\int_0^Q G''(x)\,\theta'(x)\,\mathrm{d}x$. Let $D^*(t) = D^*(0,t)$ be the demand that occurs in t days. $D^*(t)$ has a gamma distribution with rate parameter 1 and shape parameter t. Let L be the lead time. The functions $f(x)$ that are of interest can be expressed in terms of the density $f_{D(L)}(x)$, the complementary distribution function $\bar{F}_{D(L)}(x) = P(D(L) > x)$, and the partial expectation $n_{D(L)}(x) = E[(D(L) - x)^+]$. We define the partial second moment by $n^2_{D(t)}(x) = E\{\frac{1}{2}[(D(L)-x)^+]^2\}$. We use the following identities, which apply to gamma demand distributions: $n_{D^*(L)}(x) = L\bar{F}_{D^*(L+1)}(x) - x\,\bar{F}_{D^*(L)}(x)$ and $n^2_{D^*(L)}(x) = [(L^2+L)/2]\cdot\bar{F}_{D^*(L+2)}(x) - L\,x\,\bar{F}_{D^*(L+1)}(x) + (x^2/2)\,\bar{F}_{D^*(L)}(x)$.

We evaluate $\int_0^Q G(x)\,\theta'(x)\,\mathrm{d}x$ as $\int_0^Q G(x)\,\mathrm{d}x + \int_0^Q G(x)\,[\theta'(x)-1]\,\mathrm{d}x$. The first of these two integrals is evaluated in closed form using the identity $\frac{\mathrm{d}}{\mathrm{d}x} n^2_D(x) = -n_D(x)$. This identity also simplifies the computations required for Zheng's policy. Integrals of the form $\int_0^x f(x)\,[\theta'(x)-1]\,\mathrm{d}x$ are written as $\int_0^{h^{-1}(x)} f(h(y))\,[\theta'(h(y)) - 1]\,h'(y)\,\mathrm{d}y$ where

$$h(y) = \mathrm{e}^{(1-1/y)} \text{ if } y \le 1, \text{ and } \; h(y) = 1 - \ln(2-y) \text{ if } 1 \le \; y \le 2. \tag{36}$$

This is numerically advantageous because both $[\theta'(h(y)) - 1]\,h'(y)$, $0 \le y \le 2$ and the interval of integration are bounded. The integrals of $G'(x)$ and $G''(x)$ are managed similarly, using the identities $\frac{\mathrm{d}}{\mathrm{d}x} n_D(x) = -\bar{F}_D(x)$ and $\frac{\mathrm{d}}{\mathrm{d}x}\bar{F}_D(x) = -f_D(x)$.

Polynomial approximation for $\theta'(x)$

Accuracy: $\pm 0.1\%$

Algorithm

$y = (1 - \mathrm{e}^{-3h^{-1}(x)})/3$ (see (36))

$\nu = c[x,0] + c[x,1]y + c[x,2]y^2 + c[x,3]y^3$

$\theta'(x) = 1 + \nu/h'(y)$ (see (36))

Data

x Range	$c[x,0]$	$c[x,1]$	$c[x,2]$	$c[x,3]$
$0 < x \leq 3.25 \times 10^{-6}$	1	3.138513	8.15216	0
$3.25{\times}10^{-6} < x \leq 0.00027$	1.130977	-3.496022	119.1526	-615.058
$0.00027 < x \leq 0.0015$	0.653346	9.77922	-0.2834	-272.15
$0.0015 < x \leq 0.0073$	0.619633	10.5112	-4.024	-273.3
$0.0073 < x \leq 0.065$	-1.44249	58.05855	-370.532	670.977
$0.065 < x \leq 0.237$	1.142996	16.91995	-152.0991	283.902
$0.237 < x \leq 0.79$	5.650726	-41.11935	97.60782	-75.0501
$0.79 < x \leq 1.0$	3.531086	-20.14113	28.72104	0
$1.0 < x \leq 1.13$	0.2771337	-0.7399454	-0.097316	0
$1.13 < x \leq 1.26$	0.348694	-0.95677	-0.114	0
$1.26 < x \leq 1.5$	0.4634541	-1.346797	0	0
$1.5 < x \leq 2.0$	-8.801133	55.67633	-87.7363	0
$2.0 < x \leq 2.7$	1.383476	-4.13484	0	0
$2.7 < x \leq 7.0$	1.1125	0	-10	0
$7.0 < x$	0	0	0	0

Polynomial approximation for $\theta(x)$

Comment: Note that by Lemma 2, $[\theta(Q) - Q] \to E[V] = 1/2$ as $Q \to \infty$.
Accuracy: $\pm 0.1\%$
Algorithm

$y = \ln(x)$

If $x \leq 0.00125$ then $\nu = [c[x,1] - c[x,3](-y)^{c[x,0]}] \,/\, [c[x,2] - c[x,4](-y)^{c[x,0]}]$

If $0.00125 < x \leq 7.0$ then $\nu = c[x,0] + c[x,1]y + c[x,2]y^2 + c[x,3]y^3 + c[x,4]y^4$

If $7.0 < x$ then $\nu = [x + \exp(-2)]/x$

$\theta(x) = x - 1/\ln[\nu \cdot x \,/(1 + e^2 x)]$

Data

x Range	$c[x,0]$	$c[x,1]$	$c[x,2]$	$c[x,3]$	$c[x,4]$
$0 < x \leq 2 \times 10^{-7}$	-1.2839	6.6085	3.76234	-30.11	-30.11
$2 \times 10^{-7} < x \leq 0.00125$	0.5292	-5.963	-2.934	-3.157	-1.78
$0.00125 < x \leq 0.029$	1.45469	0.523461	0.180076	0.0213841	0.00088453
$0.029 < x \leq 0.15$	1.04877	0.013856	-0.058843	-0.0283627	-0.0030053
$0.15 < x \leq 0.62$	1.050675	0.024767	-0.038544	-0.014739	0
$0.62 < x \leq 3.0$	1.04993	0.02395	-0.03346	-0.003828	0.0069
$3.0 < x \leq 7.0$	1.07574	-0.031276	-0.002662	0.002	0

Polynomial approximation for $\eta(x) \equiv \int_0^x z\cdot(\theta'(z)-1)\,\mathrm{d}z$

Accuracy: $\pm 0.25\%$
Algorithm
$y = \ln(x)$
If $x \le 0.0025$ then $\nu = \left[1 - c[x,2]\cdot(c[x,4]-y)^{c[x,0]}\right] / \left[c[x,1] - c[x,3]\cdot(c[x,4]-y)^{c[x,0]}\right]$
If $0.0025 < x \le 1.4$ then $\nu = c[x,0] + c[x,1]y + c[x,2]y^2 + c[x,3]y^3 + c[x,4]y^4$
If $x \le 1.4$ then $\eta(x) = \nu\,(1-\mathrm{e}^{-x}) / \{11.93707[1-(8x+1)\mathrm{e}^{-8x}] + (\ln(\min(x,0))^2\}$
If $1.4 < x$ then $\eta(x) = c[x,4]\cdot\left\{1 - (c[x,0]\,x + c[x,1])\mathrm{e}^{-c[x,2]x}/(x+c[x,3])^2\right\}$
Data

x Range	$c[x,0]$	$c[x,1]$	$c[x,2]$	$c[x,3]$	$c[x,4]$
$0 < x \le 0.0025$	0.772	0.9062	0.2774	0.2774	1.85
$0.0025 < x \le 0.035$	-0.0103	-0.16226	0.00445	0.001533	
$0.035 < x \le 0.105$	1.7647	1.338	0.42813	0.0415	
$0.105 < x \le 0.36$	0.6827	-0.9955	-1.3737	-0.5603	-0.0741536
$0.36 < x \le 1.4$	0.9877	0.057	-0.0774	0.063	
$1.4 < x$	-2.02	17.454	0.6573	4.558	1/12

We gratefully acknowledge the helpful comments of the anonymous referees.

References

[1] S. Asxater, Using the Deterministic EOQ Formula in Stochastic Inventory Control. *Management Sci.* **42** (1996) 830-834.
[2] D. Beyer and S. Sethi, Average Cost Optimality in Inventory Models with Markovian Demands. *J. Optim. Theory Appl.* **92** (1997) 497-526.
[3] S. Bollapragada, A simple heuristic for computing nonstationary (s,S) policies. *Oper. Res.* **47** (1999) 576-585.
[4] S. Browne and P. Zipkin, Inventory Models with Continuous, Stochastic Demands. *Ann. Appl. Probab.* **1** (1991) 419-435.
[5] F. Chen and Y. Zheng, Inventory Policies with Quantized Ordering. *Naval Res. Logist.* **39** (1992) 654-665.
[6] K. Cheung, A Continuous Review Inventory Model with a Time Discount. *IEEE Trans.* **30** (1998) 747-757.
[7] A. Federgruen and Y. Zheng, Computing an Optimal (s,S) Policy is as Easy as Evaluating a Single Policy. *Oper. Res.* **39** (1991) 654-665.
[8] A. Federgruen and P. Zipkin, Computational Issues in an Infinite-Horizon, Multi-Echelon Inventory Model. *Oper. Res.* **32** (1984) 818-835.
[9] A. Federgruen and P. Zipkin, An Efficient Algorithm for Computing Optimal (s,S) Policies. *Oper. Res.* **32** (1984) 1268-1285.
[10] A. Federgruen and P. Zipkin, Computing Optimal (s,S) Policies in Inventory Models with Continuous Demands. *Adv. in Appl. Probab.* **17** (1985) 424-442.
[11] W. Feller, *An Introduction to Probability and its Applications*, Vol. II. Wiley, New York (1966).
[12] M. Fu, Sample Path Derivatives for (s,S) Inventory Systems. *Oper. Res.* **42** (1994) 351-364.

[13] J. Hu, S. Nananukul and W. Gong, A New Approach to (s, S) Inventory Systems. *J. Appl. Probab.* **30** (1993) 898-912.
[14] G. Gallego, New Bounds and Heuristics for (Q, r) Policies. *Management Sci.* **44** (1988) 219-233.
[15] G. Gallego and T. Boyaci, *Managing Waiting Time Related Service Levels in Single-Stage (Q,r) Inventory Systems*, Working paper. Department of Industrial Engineering and Operations Research, Columbia University, New York, NY (2000).
[16] G. Gallego and T. Boyaci, *Minimizing Holding and Ordering Costs subject to a Bound on Backorders is as Easy as Solving a Single Backorder Cost Model*, Working paper. Department of Industrial Engineering and Operations Research, Columbia University, New York, NY (2000).
[17] G. Gallego and T. Boyaci, *Minimizing Average Ordering and Holding Costs subject to Service Constraints*, Working paper. Department of Industrial Engineering and Operations Research, Columbia University, New York, NY (2000).
[18] G.J. Hadley and T.M. Whitin, *Analysis of Inventory Systems*. Prentice Hall, Englewood Cliffs, NJ (1963).
[19] T. Hida, *Stationary Stochastic Processes*. Princeton University Press, Princeton, NJ (1970).
[20] L.A. Johnson and D.C. Montgomery, *Operations Research in Production Planning, Scheduling, and Inventory Control*. John Wiley and Sons, New York (1974).
[21] S. Nahmias, *Production and Operations Analysis*, Second Edition. Irwin, Homewood Illinois, 60430 (1993).
[22] N.U. Prabhu, *Stochastic Storage Processes*. Springer-Verlag, New York (1980).
[23] S. Resnick, *Adventures in Stochastic Processes*. Birkhauser, Boston, MA (1992).
[24] L. Robinson, Tractible (Q, R) Heuristic Models for Constrained Service Levels. *Management Sci.* **43** (1997) 951-965.
[25] R. Roundy and G. Samorodnitsky, *Optimal and Heuristic (s,S) Inventory Policies for Levy Demand Processes*, Technical Report. School of Opeations Research and Industrial Engineering, Cornell University, Ithaca NY 14853 (1996).
[26] I. Sahin, On the Objective Function Behavior in (s, S) Inventory Models. *Oper. Res.* **82** (1982) 709-724.
[27] R. Serfozo and S. Stidham, Semi-Stationary Clearing Processes. *Stochastic Process. Appl.* **6** (1978) 165-178.
[28] M. Sharpe, *General Theory of Markov Porcesses*. Academic Press, Boston Massachusetts (1988).
[29] J. Song and P. Zipkin, Inventory Control in a Fluctuating Demand Environment. *Oper. Res.* **41** (1993) 351-370.
[30] T.E. Vollman, W.L. Berry and D.C. Whybark, *Manufacturing Planning and Control Systems*, Third Edition. Irwin, Homewood Illinois (1992).
[31] Y. Zheng and A. Federgruen, Finding Optimal (s, S) Policies is About as Simple as Evaluating a Single Policy. *Oper. Res.* **39** (1991) 654-665.
[32] Y. Zheng, On Properties of Stochastic Inventory Systems. *Management Sci.* **38** (1992) 87-103.
[33] P. Zipkin, Stochastic Lead Times in Continuous-Time Inventory Models. *Naval Res. Logist. Quarterly* **33** (1986) 763-774.
[34] P. Zipkin, *Foundations of Inventory Management*. McGraw-Hill, Boston Massachusetts (2000).

to access this journal online:
www.edpsciences.org