OPERATIONS RESEARCH
Vol. 56, No. 2, March–April 2008, pp. 471–486
ISSN 0030-364X | EISSN 1526-5463 | 08 | 5602 | 0471

DOI 10.1287/opre.1070.0442

Regulated Random Walks and the LCFS Backlog Probability: Analysis and Application

Opher Baron
Rotman School of Management, University of Toronto, Toronto, Ontario, Canada M5S 3E6,
opher.baron@rotman.utoronto.ca

Random walks have been used extensively within operations research models such as inventory systems and single-server queues to estimate performance measures. In this paper, we use sample-path analysis to express the steady-state probability of a one-sided regulated random walk to increase and be above a threshold, referred to as the last-come-first-serve (LCFS) backlog probability. We approximate the LCFS backlog probability under mild assumptions on the distribution of the random walk's steps and provide its exact expression when the steps are exponentially distributed, and a closed-form approximation when the steps are normally distributed. In our numerical experiments, the average relative gap between the approximated LCFS backlog probabilities and their simulated values is 5.13%. We further show that the LCFS backlog probability is an upper bound on the loss probability—the probability that a two-sided regulated random walk is at a boundary. This bound is tighter than the backlog probability—the probability that a random walk ever crosses a threshold—that also bounds the loss probability. In an inventory application, we demonstrate that using the LCFS backlog probability rather than the backlog probability reduces the inventory level required to satisfy a service-level constraint on the percentage of orders backlogged. In our examples, this reduction leads to cost savings of 31% on average.

Subject classifications: inventory/production, uncertainty, stochastic: base-stock level for a single item; probability, random walk: bounds on threshold-related probabilities.
Area of review: Stochastic Models.
History: Received January 2006; revisions received November 2006, February 2007; accepted March 2007. Published online in *Articles in Advance* January 14, 2008.

1. Introduction

The one-dimensional random walk and its regulated versions, the one- and two-sided regulated random walks, are used extensively in the operations research literature. For example, they are used in models of inventory processes, queues, and token-bucket admission control and pricing schemes. The last application, described in Baron (2003) and Baron et al. (2005), provides the original motivation for this research. In the management of each of these applications, one often needs to estimate certain probabilities to determine appropriate service levels or controls.

We focus on three such probabilities and name them in accordance with inventory terminology (described in §1.1.2). The first probability is the steady-state probability that a one-sided regulated random walk is above a threshold. We refer to this probability as the *backlog probability.* The second probability is the steady-state probability that a two-sided regulated random walk is on a particular boundary. We refer to this probability as the *loss probability.* The third probability we introduce is the steady-state probability that a one-sided regulated random walk increases and is above a threshold. We refer to this probability as the *last-come-first-serve* (*LCFS*) *backlog probability.* The main purpose of this paper is to present an analytic expression for the LCFS backlog probability and show that it is an upper bound on the loss probability. An additional result of interest is the moment-generating functions for both the overshoots and undershoots of a threshold by random walks involving exponential steps.

It is trivial to show that the loss probability is lower than the backlog probability. Moreover, in their inventory model, Paschalidis and Liu (2003) prove the equivalence of the exponential decay rate of both probabilities (for steps that satisfy a large deviation principle). Therefore, one might use the backlog probability as a bound on the loss probability (e.g., §2.4.1 in Whitt 2002). However, this method overestimates the loss probability. The result could be excessive inventory levels and unnecessary costs. These costs can be reduced using the loss probability or a tighter upper bound on it. Clearly, the loss probability can be calculated if the distribution of a two-sided regulated random walk were known. However, expressing this probability in closed form is still an open problem, even for random walks involving exponential steps. Therefore, it is a common practice to approximate this distribution using Brownian motion based on the first two moments of the step size (e.g., Berger and Whitt 1992, Harrison 1985, and Kushner and Dupuis 2001) or by discretizing the state space in the range of the random walk and analyzing the resulting finite-state Markov chain (Gallager 1996).

In this paper, we use a different approach: We rely on the sample paths of the regulated random walks to relate the three probabilities defined above. Using this approach, we show that the loss probability is lower than the LCFS backlog probability, which in turn is lower than the backlog probability. Thus, the LCFS backlog probability is a tighter upper bound on the loss probability than is the backlog probability. Our main contribution is an analytic expression for the LCFS backlog probability. Moreover, under mild assumptions on the random walks steps' distributions, we present a methodology to express this probability.

This methodology leads to an exact expression for the LCFS backlog probability in the case of double-exponential steps, and to approximations for shifted-exponential or normal steps (we define double- and shifted-exponential steps in §4). The exact expression for the LCFS backlog probability shows that this probability (and therefore the loss probability) might be substantially lower than the backlog probability. This finding is supported by our approximations and simulation results for the other two steps' distributions. Thus, using this paper's results may significantly decrease the error in estimating the LCFS backlog and loss probabilities. In the inventory context, we demonstrate in §5 that using our results might improve the choice of inventory levels while maintaining required service levels. We compare our results to ones by Paschalidis and Liu (2003) and show that our results can be more accurate.

This paper is organized as follows. Next, we formally define the probabilities of interest and discuss them in the contexts of inventory control, queueing, and token-bucket admission controls. Section 2 relates the backlog probability and the LCFS backlog probability and uses the latter as an upper bound on the loss probability. Section 3 describes a methodology to express the LCFS backlog probability under mild assumptions on the steps' distribution (namely, that steps follow a nonarithmetic distribution and have a conjugate point; both are defined in §3). Section 4 demonstrates this methodology for double-exponential, shifted-exponential, and normal steps. Section 5 presents a numerical study for the accuracy of our approximations and demonstrates potential savings based on our results for a simple inventory problem. We summarize the paper in §6.

1.1. Preliminaries: Probabilities of Interest and Their Applications

In this section, we formally define a one-dimensional random walk, its one- and two-sided regulated versions, and three probabilities of interest: the backlog, loss, and LCFS backlog probabilities. We then review an inventory control application of random walks that clarifies why we refer to these probabilities as such. We also briefly review the applications of random walks and these probabilities to queueing and token-bucket admission control.

1.1.1. Definitions. Let $\{X_i\}_{i=1}^{\infty}$ be a sequence of independent and identically distributed (i.i.d.) random variables, and let X be the generic random variable of this sequence. Assume that the c.d.f., $F_X(x)$, is known and satisfies $F_X(0) < 1$ and $E(X) < 0$. Then, the *random walk* $S_n = \sum_{i=1}^{n} X_i$ (with $S_0 \equiv 0$) tends to $-\infty$ with probability one (e.g., Ross 1983).

A *one-sided regulated random walk* $\{Y_i\}_{i=1}^{\infty}$ that is regulated at zero is given by

$$Y_0 = 0 \quad \text{and} \quad Y_{i+1} = \max\{0, Y_i + X_{i+1}\} \quad \text{for } i = 0, 1, \ldots. \tag{1}$$

A *two-sided regulated random walk* $\{\tilde{Y}_i\}_{i=1}^{\infty}$ that is regulated at zero and at $d > 0$ is given by

$$\tilde{Y}_0 = 0 \quad \text{and} \quad \tilde{Y}_{i+1} = \min\{d, \max\{0, \tilde{Y}_i + X_{i+1}\}\} \quad \text{for } i = 0, 1, \ldots. \tag{2}$$

Let $I\{A\}$ be the indicator function with a value of one if A happens and of zero otherwise. Then, we focus on the probabilities:

the backlog probability:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I\{Y_i > d\} = P(Y > d); \tag{3}$$

the loss probability:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I\{\tilde{Y}_i = d\} = P(\tilde{Y} = d); \quad \text{and} \tag{4}$$

the LCFS backlog probability:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I\{Y_i > d, X_i > 0\} \equiv P_{\text{LCFS}}(d). \tag{5}$$

Here, we used the random variables $Y = \lim_{i\to\infty} Y_i$ and $\tilde{Y} = \lim_{i\to\infty} \tilde{Y}_i$, where these limits are defined in the weak convergence sense (and the first exists because $E(X < 0)$). The limit $P_{\text{LCFS}}(d)$ in (5) exists because the series is nonnegative and is dominated by the convergent series in (3). We note that for $d = 0$, the loss probability and the LCFS backlog probability are identical.

When X is not continuous, we define the loss probability as

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I\{\tilde{Y}_i = d, X_i > d - \tilde{Y}_{i-1}\}.$$

With this definition for the loss probability, our analysis is easily carried to include Xs that are not continuous. However, to simplify the exposition, we assume that X is continuous, and focus on the loss probability as defined in (4).

1.1.2. Inventory and Production Application. We review an inventory model similar to the ones in Glasserman and Tayur (1996), Glasserman (1997), Bradely and Glynn (2002), and Paschalidis and Liu (2003). Consider a production facility where both demand and capacity in each period are i.i.d. random variables such that the mean demand is strictly smaller than the mean capacity. The production facility operates according to a base-stock level inventory control policy. Time is divided into periods with fixed length. Given a base-stock level, d, production in any period continues until there are d products stored or the capacity in that period is fully utilized. We denote the inventory level by I_i, the demand by D_i, and the capacity by C_i, all for the ith period. At the beginning of each period, the demand and capacity are determined; then, the production decision is made. We assume, as in Glasserman (1997), that units produced in a period are available before the period ends. Demand is satisfied at the end of each period. We consider three cases regarding how shortages are treated: unfilled demand is (i) fully backlogged and later fulfilled in a first-come-first-serve (FCFS) manner, (ii) lost, or (iii) fully backlogged and later fulfilled in an LCFS manner. We name the backlog with the FCFS model simply *the backlog model*, as is standard in the literature.

In the backlog model, the inventory at the $i+1$th period is

$$\begin{aligned} I_{i+1} &= I_i - D_{i+1} + \min\{C_{i+1}, d - I_i + D_{i+1}\} \\ &= \min\{I_i - D_{i+1} + C_{i+1}, d\}. \end{aligned} \tag{6}$$

We define the shortfall process $Y_i = d - I_i$, observe that it is nonnegative, and use (6) to write

$$\begin{aligned} Y_{i+1} &= d - \min\{I_i - D_{i+1} + C_{i+1}, d\} \\ &= \max\{Y_i + D_{i+1} - C_{i+1}, 0\}. \end{aligned}$$

Let $X_i = D_i - C_i$ and assume that period 0 starts with inventory level d. Then, the shortfall process is a one-sided regulated random walk identical to (1), and the backlog probability—i.e., the probability that orders are backlogged—is given by (3).

In the lost demand case, the inventory at the $i+1$th period is

$$\tilde{I}_{i+1} = \max\{\min\{C_{i+1} + \tilde{I}_i - D_{i+1}, d\}, 0\}.$$

An analysis similar to that of the backlog case shows that the shortfall process for this case is a two-sided regulated random walk identical to (2), and therefore the probability that orders are lost is given by (4).

The model of backlog demand that is fulfilled in an LCFS manner might be interesting in its own right. Assume that the backlog's cost has two components: the first is a one-time cost associated with loss of goodwill and the second is a cost proportional to the backlog's size and length. Once shortages occur, we can allocate production in an LCFS manner, i.e., to first satisfy new orders such that they would not be backlogged, and then to satisfy as much of the backlogged demand as possible. Such an allocation decreases the number of periods where some of the new demand is backlogged; however, it leaves the inventory level the same as in the standard backlog model. Thus, the LCFS allocation decreases the cost associated with loss of goodwill without increasing costs associated with the backlog size and duration.

To express the probability that new demand is backlogged when allocating production in an LCFS manner (the LCFS backlog probability), observe that if the demand in the ith period is lower than the capacity, the capacity would satisfy all of the period's demand. Therefore, the demand in the ith period is backlogged only if the total demand in this period is higher than the capacity, and at the end of this period the shortfall is greater than d (i.e., there is a backlog at the end of the ith period). Mathematically, we will have an LCFS backlog in the ith period only if $I\{Y_i > d, X_i > 0\} = 1$. Therefore, the probability that orders are backlogged when backlog is allocated in an LCFS manner is given by (5).

1.1.3. Queueing. It is well known (e.g., Cohen 1982) that the Lindley recursion describing the waiting time in a GI/G/1 queue, with the FCFS service discipline, is equivalent to (1). Thus, the backlog probability in (3) describes the steady-state probability of waiting d time units or more in such a queue.

Restricted-access queues are used to analyze finite dams in Prabhu (1980) and are discussed in Cohen (1982, Part III). We consider a restricted-access queue with deterministic interarrival times of length r and general i.i.d. service times where the workload (or virtual waiting time) is restricted to be lower than $d+r$. Thus, a job whose admission would increase the workload beyond $d+r$ is only partially accepted to increase the system's workload to $d+r$. Similar to the Lindley recursion, the waiting time in such a restricted-access queue is equivalent to (2) (e.g., Cohen 1982, p. 508). Therefore, the loss probability in (4) describes the probability of customers waiting d in such a restricted-access queue. This probability is also equivalent to the proportion of jobs that are only partially processed by the server.

1.1.4. Token-Bucket Admission Controls. Two common admission controls in the network and telephony literature are the *token bucket with rate control* (Berger 1991) and the *token bucket* (Berger and Whitt 1992). They use two parameters to regulate the demand for a network's resources: a token rate, denoted r, and a bucket depth, denoted d. Every demand source gets tokens at a rate r (not necessarily an integer) and spends a token to send a packet of data through the network. Tokens that are not used can be accumulated in a bucket of size d, and if a token arrives when the bucket is full, it is lost. We assume that demand in each period is i.i.d. with a mean that is smaller than r.

In the token bucket with rate control case, when there are no tokens in the bucket, packets are backlogged and the number of tokens in the bucket evolves as in (6) (with $C_i = r$). Thus, the bucket level is a one-sided regulated random walk, as given in (1), and the percentage of periods with backlog of packets is identical to the backlog probability in (3). In the token-bucket case, when there are no tokens in the bucket, packets are lost and the number of tokens in the bucket evolves as in (2). Thus, the bucket level is a two-sided regulated random walk and the percentage of periods with losses of packets is identical to the loss probability in (4). Further discussion of the token-bucket admission controls and their application to pricing can be found in Baron et al. (2005) and Baron (2003).

Here, as well as in the restricted-access queue discussed above, the LCFS backlog probability can serve as a bound on the loss probability.

2. The LCFS Backlog Probability as an Upper Bound on the Loss Probability

To establish sample path relations between the one- and two-sided regulated random walks, we define cycles and subcycles. To ease the exposition, we depict, in Figure 1, a cycle and subcycles for a realization of one- and two-sided regulated random walks. Using these definitions, we relate the probabilities of interest in Theorem 1. This theorem expresses the LCFS backlog probability as a constant (smaller than one) times the backlog probability. This constant is closely related to the expected length of a subcycle seen by a random observer, which we call a *representative subcycle*. Therefore, to make Theorem 1 useful, we will devote the next section to expressing the expected length of a representative subcycle.

2.1. Preliminaries: Cycles and Subcycles

Definition 1 (Cycles). Let $\Psi^0 = 0$. For $j \geqslant 1$, define $W^j \equiv \inf\{i \geqslant \Psi^{j-1}\colon Y_i > d\}$ and $\Psi^j \equiv \inf\{i \geqslant W^j\colon Y_i = 0\}$. We call the interval $[\Psi^{j-1}, \Psi^j)$ the jth cycle and its length is $\Theta_j = \Psi^j - \Psi^{j-1}$.

The sequence $W^1, \Psi^1, W^2, \Psi^2, \ldots$ lists periods in which the "other" boundary is hit for the first time since the last hit of a boundary. It is a sequence of stopping times with respect to the process $\{Y_i\}$. Then, a cycle is the time interval between two events of hitting zero if the last boundary hit was at d. Note that the beginning of each cycle defines a renewal process, and therefore cycles are i.i.d. We denote the expected length of a cycle by $E(\Theta)$ and note that $E(\Theta) < \infty$ (because $E(X) < 0$). In Figure 1, we have $W^1 = 6$, $\Psi^1 = 28$, and $W^2 > 35$; thus, $\Theta_1 = 28$ and the second cycle does not end in the figure.

Figure 1. One and a "half" cycles for negative-drift one-sided (with stars) and two-sided (with circles) regulated random walks.

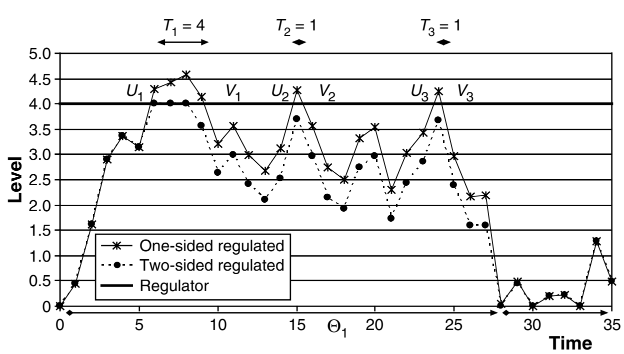

Note. Both are regulated at zero; the two-sided regulated walk is also regulated at $d = 4$.

For the jth cycle we define:

Definition 2 (Subcycles). Let $U_1^j \equiv \inf\{i \geqslant \Psi^{j-1}\colon Y_i > d\}$. For $k \geqslant 1$, define $V_k^j \equiv \inf\{i \geqslant U_k^j\colon Y_i \leqslant d\}$, and for $k \geqslant 2$, define $U_k^j \equiv \inf\{i \geqslant V_{k-1}^j\colon \{Y_i > d\} \cup \{Y_i = 0\}\}$. For $k \geqslant 1$, let $T_k^j \equiv V_k^j - U_k^j$. Let $K^j \equiv \max\{k \geqslant 1\colon Y_{U_k^j} > 0\}$; then, for $k = 1, \ldots, K^j$, we call the interval $[U_k^j, V_k^j)$ the kth subcycle of the jth cycle, and its length is given by T_k^j.

Note that $U_1^j = W^j$ by Definition 1 and K^j is the number of subcycles in the jth cycle. Moreover, once $Y_i = 0$ for $i > W^j$, the cycle ends; thus, we have $U_k^j = V_k^j = \Psi^j$ and $T_k^j = 0$ for $k > K^j$. By definition, $E(T_k^j \mid K^j \geqslant k) \geqslant 1$; yet $\sum_{k=1}^{\infty} E(T_k^j) < \infty$ because $\sum_{k=1}^{\infty} T_k^j = \sum_{k=1}^{K^j} T_k^j \leqslant \Theta_j$ and $E(\Theta) < \infty$. Similarly, $E(K^j) < \infty$, and because $K^j \geqslant 1$ and is nonnegative, $E(K^j) = \sum_{k=1}^{\infty} P(K^j \geqslant k)$. In the sequel, we focus on a single cycle and omit the superscript j when no confusion arises. In Figure 1, we have $U_1 = 6$, $U_2 = 15$, $U_3 = 24$, $V_1 = 10$, $V_2 = 16$, $V_3 = 25$, and $U_k = V_k = 28$ for $k \geqslant 4$; thus, $T_1 = 4$, $T_2 = T_3 = 1$, $K = 3$, and $T_k = 0$ for $k \geqslant 4$.

We emphasize that, in general, lengths of different subcycles are not i.i.d.: A subcycle's length depends on the overshoot at the beginning of the subcycle. These overshoots are not necessarily independent or identically distributed between different subcycles. Observe that if X_i has a positive exponential component, i.e., $X_i = D_i - C_i$ where $D_i \sim \exp(\mu)$, C_i are i.i.d., and $E(C) > 1/\mu$ to satisfy $E(X) < 0$, the overshoots of subcycles are identically distributed due to the memoryless property. Thus, in this case, subcycles' lengths are i.i.d.

2.2. Relating the Probabilities of Interest

It follows from the definitions of Y_i and $\tilde{Y}_i$ that $\tilde{Y}_i \leqslant Y_i$. An example is given in Figure 1; see also Paschalidis and Liu (2003). Therefore, $F_Y(y) \leqslant F_{\tilde{Y}}(y)$ and

$$P(\tilde{Y} = d) \leqslant P_{\mathrm{LCFS}}(d) \leqslant P(Y > d). \tag{7}$$

The right inequality is clear because the indicator in (3) is taken over a less restricted set than in (5). To establish the left inequality, observe that $I\{\tilde{Y}_i = d\} = I\{\tilde{Y}_i = d, X_i > 0\}$ (because X is continuous). Moreover, because $\tilde{Y}_i \leqslant Y_i$, we have $I\{\tilde{Y}_i = d, X_i > 0\} \leqslant I\{Y_i \geqslant d, X_i > 0\}$; thus, because the event $\{Y_i \geqslant d\}$ is identical to the event $\{Y_i > d\}$ almost surely, the inequality follows from the definitions of $P(\tilde{Y} = d)$ and $P_{\mathrm{LCFS}}(d)$.

In the inventory application, any period with loss is also a period with backlog and has $X_i > 0$, that is, this period has backlog in the LCFS model as well. However, there might be periods with $X_i > 0$ in which there are backlogs at the LCFS model but no losses. In such periods, $\widetilde{Y}_i < d$, but both $Y_i > d$ and $X_i > 0$. In Figure 1, such periods are 15 and 24 at the second and third upcrossing of d by the one-sided regulated random walk.

Using the renewal reward theorem and the definitions of cycles and subcycles:

$$P(Y > d) = \frac{E(\text{number of periods with backlog in a cycle})}{E(\text{number of periods in a cycle})} = \frac{E\left(\sum_{j=1}^{K} T_j\right)}{E(\Theta)} \tag{8}$$

and

$$P_{\text{LCFS}}(d) = \frac{E(\text{number of periods with backlog and } X_i > 0 \text{ in a cycle})}{E(\text{number of periods in a cycle})} = \frac{m}{E(\Theta)}, \tag{9}$$

where m denotes the expected number of periods with backlogs and $X_i > 0$ in a cycle.

Before expressing m, we review Wald's equality (e.g., Siegmund 1985). Wald's equality states that for a stopping time N with $E(N) < \infty$ and an i.i.d. random variable X_i with $E[X_i] = E(X)$,

$$E(X) = \frac{E\left[\sum_{i=1}^{N} X_i\right]}{E(N)}. \tag{10}$$

Let $\bar{F}_X(X) \equiv 1 - F_X(X)$ be the complement c.d.f. of X:

Lemma 1. *The expected number of periods with backlogs and $X_i > 0$ in a cycle, m, is*

$$m = \bar{F}_X(0) E\left(\sum_{j=1}^{K} T_j\right) + E(K).$$

Proof. By the definition of U_j, V_j, and K, the number of periods within subcycles in a cycle (i.e., the number of periods with backlog in a cycle) is $\sum_{j=1..K} \sum_{U_j \leqslant i < V_j} 1$. Moreover, the probability that the ith step is positive equals $I\{X_i > 0\}$. Thus,

$$m = E\left[\sum_{j=1}^{K}\left(\sum_{i=u_j}^{V_j-1} I\{X_i > 0\}\right)\right] = \sum_{k=1}^{\infty}\left[\sum_{j=1}^{k} E\left(\sum_{i=u_j}^{V_j-1} I\{X_i > 0\}\right)\right] P(K = k).$$

To calculate the expectation, we observe that for any jth subcycle,

$$\sum_{i=U_j}^{V_j-1} I\{X_i > 0\} = 1 + \sum_{i=U_j+1}^{V_j} I\{X_i > 0\} \tag{11}$$

because $I\{X_{U_j} > 0\} = 1$ and $I\{X_{V_j} > 0\} = 0$ in subcycles.

Observe that by the strong Markov property for random walks (that is, that a random walk starts afresh at each stopping time) and because U_1 is a nondefective stopping time, we can apply Wald's equality starting at U_1. Thus, we consider U_j and V_j as stopping times and $I\{X_i > 0\}$ the i.i.d random variable, with expectation $\bar{F}_X(0)$. Then, for any jth subcycle, the expectation of the summation in the right-hand side of (11) is

$$E\left[\sum_{i=U_1}^{V_j} I\{X_i > 0\}\right] - E\left[\sum_{i=U_1}^{U_j} I\{X_i > 0\}\right] = \bar{F}_X(0)[E(V_j) - E(U_j)] = \bar{F}_X(0) E(T_j \mid K \geqslant j).$$

Therefore,

$$E\left(\sum_{i=U_j}^{V_j-1} I\{X_i > 0\}\right) = \bar{F}_X(0) E(T_j \mid K \geqslant j) + 1$$

and

$$m = \sum_{k=1}^{\infty}\sum_{j=1}^{k}[(\bar{F}_X(0) E(T_j \mid K \geqslant j)) + 1] P(K = k) = \sum_{j=1}^{\infty}[(\bar{F}_X(0) E(T_j \mid K \geqslant j)) + 1] P(K \geqslant j). \tag{12}$$

Observe that for any T_j (the length of a jth subcycle),

$$E(T_j) = E(T_j \mid K \geqslant j) P(K \geqslant j) + E(T_j \mid K < j) P(K < j) \tag{13}$$

$$= E(T_j \mid K \geqslant j) P(K \geqslant j) \tag{14}$$

$$= E(T_j I\{K \geqslant j\} \mid K \geqslant j) P(K \geqslant j)$$

$$= E(T_j I\{K \geqslant j\}), \tag{15}$$

where (13) follows from the total expectation theorem, (14) holds because given $K < j$, $E(T_j) = 0$, and the third equality holds because given $K \geqslant j$, $I\{K \geqslant j\} = 1$. Applying similar logic to $T_j I(K \geqslant j)$ establishes (15).

Thus, from (14) and (15),

$$\sum_{j=1}^{\infty} E(T_j \mid K \geqslant j) P(K \geqslant j) = \sum_{j=1}^{\infty} E(T_j I\{K \geqslant j\}) = E\left(\sum_{j=1}^{K} T_j\right), \tag{16}$$

where the last equality follows by definition. Substituting (16) and $\sum_{j=1}^{\infty} P(K \geqslant j) = E(K)$ into (12) and using some algebra concludes the proof. □

Using Lemma 1, we provide in Theorem 1 below an analytic bound on the loss probability. In contrast, Paschalidis and Liu (2003) suggested a nonanalytic approximation for the loss probability based on simulation and their analytic results.

Theorem 1. *The loss probability in a two-sided regulated random walk, $P(\tilde{Y}=d)$, the LCFS backlog probability $P_{LCFS}(d)$, and the backlog probability in a one-sided regulated random walk, $P(Y>d)$, are related via*

$$P(\tilde{Y}=d) \leqslant P_{LCFS}(d) = cP(Y>d), \tag{17}$$

where

$$c \equiv \left(\bar{F}_X(0) + \frac{E(K)}{E\left(\sum_{j=1}^{K} T_j\right)}\right) < 1. \tag{18}$$

Proof. The left inequality of (17) is from (7); thus, we focus on establishing the equality in (17) and the inequality in (18). Substituting m from Lemma 1 into (9),

$$P_{\text{LCFS}}(d) = \frac{\bar{F}_X(0)E\left(\sum_{j=1}^{K} T_j\right) + E(K)}{E(\Theta)}.$$

Using (8), we can replace $E(\Theta)$ with $E(\sum_{j=1}^{K} T_j)/P(Y>d)$ to prove (17).

To establish (18), we think of a subcycle as a series of trials to end it. Because the overshoot of d before each trial is positive, the probability of ending the subcycle at each trial is strictly smaller than $F_X(0)$. Thus, for each subcycle, $E(T_j)$ is strictly larger than $1/F_X(0)$—the expected number of trials of a geometric random variable with success probability $F_X(0)$. From this,

$$E\left(\sum_{j=1}^{K} T_j\right) > \frac{E\left(\sum_{j=1}^{K} 1\right)}{F_X(0)} = \frac{E(K)}{F_X(0)}, \tag{19}$$

and using some algebra concludes the proof. □

To use Theorem 1, we still need to express $E(K)/E(\sum_{j=1}^{K} T_j)$. To this end, we define a *representative subcycle* and denote its length by R. We show that $E(R) = E(\sum_{j=1}^{K} T_j)/E(K)$. Consider the sequence of K^i, $i = 1, 2, \ldots,$ representing the number of subcycles in the ith cycle from Definition 2. Recall that the K^i are discrete i.i.d. random variables. Consider a discrete time index $t = 1, 2, \ldots,$ and define

$$H(t) = t - \sup\left\{\sum_{i=1}^{l} K^i \colon \sum_{i=1}^{l} K^i < t \text{ and } l = 1, 2, \ldots\right\}. \tag{20}$$

When counting only subcycles, the second term in (20) provides the time index when the last cycle before t ended. Let l^* be the upper limit of the summation in this term (thus, $H(t) \in 1, \ldots, K^{l^*+1}$). Then, $H(t) = k$ implies that the tth subcycle in the process is the kth subcycle in the l^*+1st cycle. Consider an observer who randomly picks a value $\check{T} \in N$. Thus, when counting only subcycles from the start of the processes, the representative subcycle is given by the $\check{T}$th subcycle. We observe that the probability that the representative subcycle is the jth subcycle in a cycle is

$$P(H(\check{T})=j) = \frac{P(K \geqslant j)}{\sum_{k=1}^{\infty} P(K \geqslant j)} = \frac{P(K \geqslant j)}{E(K)}. \tag{21}$$

Using (21), a representative subcycle can be viewed as a convex combination of subcycles. We observe that when a jth subcycle is chosen as the representative subcycle, its expected length is $E(T_j \mid K \geqslant j)$. Thus, using the total expectation theorem,

$$E(R) = \sum_{j=1}^{\infty} E(T_j \mid K \geqslant j) P(H(\check{T})=j). \tag{22}$$

Note that, in general, the distribution of the length of a representative subcycle differs from the distributions of lengths of actual subcycles. We next express the expected length of a representative subcycle:

Theorem 2. *The expected length of a representative subcycle is*

$$E(R) = \frac{E\left(\sum_{j=1}^{K} T_j\right)}{E(K)}.$$

Proof. Using (21) and (22)

$$E(R) = \sum_{j=1}^{\infty} \frac{E(T_j \mid K \geqslant j) P(K \geqslant j)}{E(K)} = \frac{E\left(\sum_{j=1}^{K} T_j\right)}{E(K)},$$

where the last equality holds by (16). □

Theorem 2 is similar in form to Wald's equality (10). However, in Theorem 2, T_j are not i.i.d. and K is not a stopping time. We attribute this similarity to the way R is defined.

We observe that c from (18) in Theorem 1 can be expressed as

$$c = \bar{F}_X(0) + \frac{1}{E(R)}.$$

Thus, using Theorem 1 requires the evaluation of $E(R)$ that for general steps' distribution is still an open problem. Therefore, we devote the next section to evaluating $E(R)$ under some minor assumptions.

3. The Expected Length of a Representative Subcycle

We use the fact that the representative subcycle is a convex combination of subcycles, and the results of Chapter 8 of Siegmund (1985), to provide a methodology to approximate $E(R)$ under mild assumptions on the steps' distribution. We note that this methodology depends only on the distribution of X; thus, it is independent of the threshold. We further discuss this independence and the validity of our approximations at the end of this section.

The kth subcycle starts with an overshoot $Y_{U_k} - d = A_k > 0$, and ends with an undershoot $d - Y_{V_k} = B_k$. In the first subcycle of Figure 1, $A_1 \approx 0.3$ (for $U_1 = 6$) and $B_1 \approx 0.8$ (for $V_1 = 10$). Thus, focusing on a single subcycle, we consider its length T as a stopping time of the negative-drift random walk with steps X_i, starting at A and downcrossing zero. Numbering periods as if the subcycle's first period is period 1, we have $Y_1 = d + A$.

Let $S^T = A + B$, i.e., for any subcycle, S^T is the total decrease of the random walk during this subcycle. Then, applying Wald's equality (10), where N is replaced by T,

$$E(T) = -E(S^T)/E(X) = \frac{E(A)+E(B)}{|E(X)|}, \tag{23}$$

where the last equality follows from linearity of the expectation operator.

3.1. Bounds on Subcycles' Expected Length

For any subcycle, we have, similar to Ross (1974):

$$\inf_{0\leqslant u} E(X_1 - u \mid X_1 > u) \leqslant E(A) \leqslant \sup_{0\leqslant u} E(X_1 - u \mid X_1 > u) \quad \text{and} \tag{24}$$

$$-\sup_{u\leqslant 0} E(X_1 - u \mid X_1 < u) \leqslant E(B) \leqslant -\inf_{u\leqslant 0} E(X_1 - u \mid X_1 < u). \tag{25}$$

Combining these with (19), (23), and that R is a convex combination of subcycles:

$$E(R) \geqslant \max\left\{\frac{1}{F_X(0)}, \frac{\inf_{0\leqslant u} E(X_1-u \mid X_1>u) - \sup_{0\leqslant u} E(X_1-u \mid X_1<u)}{|E(X)|}\right\}, \tag{26}$$

$$E(R) \leqslant \frac{\sup_{0\leqslant u} E(X_1-u \mid X_1>u) - \inf_{u\leqslant 0} E(X_1-u \mid X_1<u)}{|E(X)|}. \tag{27}$$

Observe that if the right-hand side of (27) is close to $1/F_X(0)$, then $P_{\mathrm{LCFS}}(d)$ and $P(Y > d)$ are close. Moreover, to be conservative in estimating $P_{\mathrm{LCFS}}(d)$, one can use the right-hand side of (27) to bound $E(R)$. Tighter expressions for the sup or inf above can be given for special distributions of X (e.g., when the positive part of X is new or better than used) similar to Ross (1974) and Glasserman (1997); thus, we do not reproduce them here.

3.2. Approximations of Subcycles' Expected Length

For the random walk S_n, we define:

Definition 3. Let $\tau^0 = 0$ and for $k \geqslant 1$ define $\tau^k \equiv \inf\{n > \tau^{k-1}\colon S_n > S_{\tau^{k-1}}\}$. Let $\tau_-^0 = 0$, and for $k \geqslant 1$ define $\tau_-^k = \inf\{n\colon n > \tau_-^{k-1},\ S_n \leqslant S_{\tau_-^{k-1}}\}$.

Then, τ^k and S_{τ^k} are the strict ladder epoch and ladder height of the random walk S_n, respectively. Note that $\tau^k - \tau^{k-1}$ and $S_{\tau^k} - S_{\tau^{k-1}}$ for $k \geqslant 1$ are (possibly degenerate, i.e., $\tau^k = \infty$ with a positive probability) i.i.d. random variables. If $S_{\tau^k} < \infty$, we let $L_\tau \overset{d}{=} S_{\tau^k} - S_{\tau^{k-1}}$, where $\overset{d}{=}$ denotes equality in distribution. Thus, L_τ is a generic ladder step, with, in general, unknown distribution. Similar to L_τ, we denote by $L_{\tau_-} \overset{d}{=} S_{\tau_-^k} - S_{\tau_-^{k-1}}$ a generic step down a ladder.

We let $G_X(s) \equiv E_X(\exp(sX))$ be the moment-generating function of X and assume that there is some s_0 such that $1 < G_X(s_0) < \infty$. Thus, s^*, the strictly positive root of

$$G_X(s^*) = 1 \tag{28}$$

exists, is unique (due to convexity of $G_X(s)$), and is called the conjugate point of X. The conjugate point exists for many commonly used distributions, such as normal. We now define the conjugate distribution (e.g., (38) of Glasserman 1997)

$$F_X^{s^*}(x) = \int_{-\infty}^{x} \frac{e^{s^* x}}{G_X(s^*)}\, dF_X(x) = \int_{-\infty}^{x} e^{s^* x}\, dF_X(x),$$

where the last equality and that $F_X^{s^*}(x)$ is a well-defined distribution function follow from (28). According to Siegmund (1985), the conjugate distribution plays the same role in estimating the overshoot as the original distribution does in estimating the undershoot.

Recall that a nonarithmetic distribution is one for which all possible values are not integer multiples of some nonnegative value. In Chapter 8 of Siegmund (1985), it is shown that for such distributions that have a conjugate point, as the threshold increases to infinity, its overshoot is distributed as the residual life of L_τ. Thus, for large d, the overshoot of the first subcycle ($A_1 = Y_{U_1} - d$) is distributed as the time to the next renewal in a renewal process with interarrival times distributed as generic ladder steps L_τs. Thus, using (8.50) of Siegmund (1985), we approximate

$$E(A_1) \approx \frac{E((L_\tau)^2)}{2E(L_\tau)} = \frac{E_{s^*}(X^2)}{2E_{s^*}(X)} - \sum_{n=1}^{\infty} \frac{E_{s^*}(S_n^-)}{n}, \tag{29}$$

where $x^- \equiv -\min\{x, 0\}$ and the expectation in (29) is taken with respect to the conjugate distribution denoted by $E_{s^*}(\cdot)$.

In a similar way, as the threshold decreases to minus infinity, its undershoot is distributed as the residual life of L_{τ_-}, which, similar to (29), is given by

$$\frac{E((L_{\tau_-})^2)}{2|E(L_{\tau_-})|} = \frac{E(X^2)}{2|E(X)|} - \sum_{n=1}^{\infty} \frac{E(S_n^+)}{n}, \tag{30}$$

where $x^+ \equiv \max\{x, 0\}$. Methods to evaluate the summations in (29) and (30) are discussed in Siegmund (1985) and Woodroofe (1979).

We suggest using the right-hand sides of (29) and (30) to approximate $E(A_k)$ and $E(B_k)$ for each k, respectively. To support this heuristic, we refer to Glasserman (1997) and Abate et al. (1995). In Glasserman (1997), it is shown that when a conjugate point exists, a change-of-measure argument can be used to express

$$P(Y > d) = \exp(-s^* d) E_{s^*}(\exp(-s^* A_1)) \quad \forall\, d > 0,$$

where, in general, the distribution of A_1 depends on d. In Abate et al. (1995), it is shown that both a heavy-traffic limit theorem and a GI/M/s model propose that for any d, the value of $E_{s^*}(\exp(-s^*A_1))$ is close to the constant $\lim_{d\to\infty} E_{s^*}(\exp(-s^*A_1))$. (This value is a constant by applying the renewal theorem; see, e.g., Chapter 8 of Siegmund 1985.) Thus, Glasserman (1997) and Abate et al. (1995) suggest that the moment-generating function of the overshoot of d evaluated at the conjugate point is similar for any d. Our heuristic borrows this idea, suggesting the use of the expected length of the overshoot of $d \to \infty$ as an approximation for the expected overshoot at the beginning of any subcycle (and similarly, the undershoot as $d \to -\infty$).

Using (21), R is a convex combination of T_k's. Moreover, we approximate $E(T_k)$ using (23) for any k. Thus, substituting (29) and (30) into (23), we have

$$E(R) \approx \left(\frac{E((L_\tau)^2)}{2E(L_\tau)} + \frac{E((L_{\tau_-})^2)}{2|E(L_{\tau_-})|}\right)\frac{1}{|E(X)|}. \qquad (31)$$

Discussion. We first observe that the approximation (31) lies between the bounds in (26) and (27) because both (29) and (30) are accurate when the absolute value of the thresholds goes to infinity. Therefore, if these bounds are close to each other, the approximation of $E(R)$ based on (31) is good. An important example is when X has an exponential component. Then, (24) and/or (25) are tight, and we expect (31) to be good. (In the following sections, we analyze two such cases.)

As noted earlier, our bounds and approximation for $E(R)$ (in (26), (27), and (31), respectively) are independent of d. Because R is a convex combination of subcycles, its length depends on the overshoots of d at the beginning of the subcycle. However, because the random walk starts afresh at the first overshoot, all these overshoots depend on d mostly through this overshoot. Moreover, Chapter 8 of Siegmund (1985) shows that for large d, the first overshoot is distributed as the residual life of L_τ, which is independent of d. Therefore, for large d, approximating $E(R)$ as if it is independent of d is reasonable.

There are three important factors that influence the accuracy of (31): the threshold d, the drift $|E(X)|$, and the standard deviation of X, denoted by σ.

We recall that the approximations used to derive (31) assume that the thresholds to be crossed are large. Thus, approximating the expected first overshoot, $E(A_1)$, using (29) improves as d increases. Unfortunately, for the expected overshoots, $E(A_j)$, for subcycles $j \geqslant 2$, the relevant thresholds are smaller than d. Thus, using (29) to approximate these $E(A_j)$ is less accurate. However, subcycles j for $j \geqslant 2$ exist with lower probability than the first one. Therefore, the approximation in (29) improves with d. Because of this improvement and because the assumption that d is independent of $E(R)$ is more reasonable as d increases, we expect the approximation in (31) to be good for large d.

As for the effect of the drift, on one hand the probability that subcycles $j \geqslant 2$ exist decreases with the drift. Thus, the error in estimating $E(A_j)$ for $j \geqslant 2$ is less significant when the drift increases. On the other hand, de Kok (1985) suggests that similar approximations to ours for the expected undershoots of a threshold ($E(B_j)$ for $j \geqslant 1$) require the threshold to be larger than the drift (but might work well even when the latter does not hold). Thus, the error in estimating $E(B_j)$ for $j \geqslant 1$ might increase with the drift. In view of the above, the effect of the drift on the accuracy of (31) is inconclusive.

We expect that the sizes of the overshoots and undershoots increase with σ, and therefore that the approximation in (31) will improve with σ.

Finally, our numerical results support the three relations discussed above.

4. Examples for Evaluating the LCFS Backlog Probability

In this section, we use Theorem 1 in conjunction with (31) to express the LCFS backlog probability in cases in which the random walk's steps follow (i) a double-exponential distribution—i.e., $X_i = D_i - C_i$ and D_i and C_i are exponentially distributed with parameters μ and λ, respectively; (ii) a shifted-exponential distribution—i.e., $X_i = D_i - r$, where D_i is exponentially distributed with parameter μ; and (iii) a normal distribution.

We present the exact closed-form expression for $E(R)$ and the LCFS backlog probability for random walks with double-exponential steps, and then use (31) to approximate this probability for random walks with shifted-exponential or normal steps. For random walks with shifted-exponential steps, the expected undershoot is approximated using the moment-generating functions of L_{τ_-} and L_τ presented in Proposition 2.

We recall that Theorem 1 states that

$$P(\tilde{Y} = d) \leqslant P_{\text{LCFS}}(d) = cP(Y > d);$$

therefore, to use Theorem 1 we need to express the backlog probability and its coefficient, given in (18). Because analytic expressions for the backlog probability are available in the literature for many cases, e.g., Glasserman (1997) and Ross (1983) for $d \to \infty$, we focus on estimating c.

4.1. Double-Exponential Steps

When the X_is follow a double-exponential distribution, we require $\lambda < \mu$ to keep $E(X) < 0$, and the LCFS backlog probability can be found in closed form:

Proposition 1. *For random walks with double-exponential steps, we have*

$$P_{LCFS}(d) = \frac{\mu}{\mu+\lambda} P(Y > d)$$
$$= \frac{\mu}{\mu+\lambda}\frac{\lambda}{\mu}\exp(-(\mu-\lambda)d). \qquad (32)$$

Proof. In this case, it is known that the backlog probability is

$$P(Y>d)=\frac{\lambda}{\mu}\exp(-(\mu-\lambda)d) \tag{33}$$

(as the waiting time in an M/M/1 queue). Thus, we need to show that $c=\mu/(\mu+\lambda)$. We observe that due to memoryless, both the left-hand and right-hand sides of (24) equal $1/\mu$ and both the left-hand and right-hand sides of (25) equal $1/\lambda$. Thus, subcycles are i.i.d. and

$$E(R)=E(T)=\left(\frac{1}{\lambda}+\frac{1}{\mu}\right)\frac{1}{|E(X)|}$$
$$=\left(\frac{\mu+\lambda}{\lambda\mu}\right)\left(\frac{1}{\lambda}-\frac{1}{\mu}\right)^{-1}=\left(\frac{\lambda+\mu}{\mu-\lambda}\right),$$

which after substituting to (18), together with $P(X>0)=\lambda/(\mu+\lambda)$, concludes the proof. □

Proposition 1 shows that for a random walk with double-exponential steps, the ratio of the backlog probability to the LCFS backlog probability approaches 2 as $\lambda\to\mu$. Thus, in highly utilized systems, using $P(Y>d)$ as a bound on $P_{\mathrm{LCFS}}(d)$ (and the loss probability) is too conservative.

4.2. Shifted-Exponential Steps

When the X_is follow shifted-exponential steps, we require $r>1/\mu$, and the c.d.f. and moment-generating function of X are

$$F_X(x)=1-\exp(-\mu(x+r))\quad \text{for } x\in[-r,\infty),$$
$$G_X(s)=\mu\frac{e^{sr}}{s+\mu}. \tag{34}$$

Again, the first overshoot of subcycles are exponentially distributed; thus, subcycles are i.i.d. and $E(R)=E(T)$. We first give the moment-generating functions of L_τ and $L_{\tau\_}$:

Proposition 2. *For random walks with shifted-exponential steps, we have*

$$G_{L_\tau}(s;\tau<\infty)=\frac{\mu}{\mu-s}\left(1-\frac{s^*}{\mu}\right)=\frac{\mu e^{-s^*r}}{\mu-s},$$
$$G_{L_{\tau\_}}(s;\tau\_<\infty)=G_{L_{\tau\_}}(s)=\mu\frac{e^{-s^*r}-e^{-sr}}{s+e^{-s^*r}\mu-\mu},$$

where s^* is the conjugate point.

Proof. The expression for $G_{L_\tau}(s;\tau<\infty)=G_{L_\tau}(s\mid\tau<\infty)P(\tau<\infty)$ follows from the fact that $G_{L_\tau}(s\mid\tau<\infty)=\mu/(\mu-s)$ based on the memoryless property, and that

$$P(\tau<\infty)=\frac{\mu-s^*}{\mu}=e^{-s^*r}$$

(e.g., (14) in Glasserman 1997), where the last equality follows from (28). The expression for the moment-generating function of $L_{\tau\_}$ follows from the expressions for $G_{L_\tau}(s;\tau<\infty)$ and $G_X(s)$ (in (34)) and the Wiener-Hopf factorization identity (e.g., Theorem 8.41 in Siegmund 1985, Asmussen 2003, p. 228):

$$G_{L_{\tau\_}}(s;\tau\_<\infty)=1-\frac{1-G_X(s)}{1-G_{L_\tau}(s;\tau<\infty)}. \quad \square$$

Observe that the methodology used to prove Proposition 2 can be generalized to additional Xs involving exponential or Erlang steps, e.g., general demand and exponential production time or general production time and exponential demand (where $P(\tau<\infty)$ is the machine's utilization due to PASTA).

Calculating the moments of $L_{\tau\_}$ based on Proposition 2, using (31), and because

$$P(Y>d)=\frac{(\mu-s^*)e^{-s^*d}}{\mu} \tag{35}$$

leads to the approximation (for random walks with shifted-exponential steps)

$$P_{\mathrm{LCFS}}(d)\approx\left(e^{-\mu r}+\left(\frac{r^2\mu^2}{2(\mu r-1)^2}+\frac{s^*-\mu}{s^*(\mu r-1)}\right)^{-1}\right)\cdot\frac{(\mu-s^*)e^{-s^*d}}{\mu}. \tag{36}$$

Observe that if our approximation is such that $c\geqslant 1$, (36) is not helpful. In such cases, we should still use the backlog probability as a bound on the LCFS and loss probabilities. (We did not encounter such cases in our numerical experiments.)

4.3. Normal Steps

When the X_is follow normal(μ,σ^2) steps, we require $\mu<0$. To estimate the expected undershoot $E(B)$, we evaluate the elements in (30) using $S_n\sim$ normal$(n\mu,n\sigma^2)$:

$$\sum_{n=1}^{\infty}\frac{E(S_n^+)}{n}=\sum_{n=1}^{\infty}\frac{1}{n\sigma\sqrt{2\pi n}}\int_0^\infty ue^{-(u-n\mu)^2/(2n\sigma^2)}\,du,$$
$$\frac{E(X^2)}{2|E(X)|}=\frac{\sigma^2+\mu^2}{-2\mu}. \tag{37}$$

Moreover, the conjugate point is $s^*=-2\mu/\sigma^2$, thus, the conjugate distribution is normal$(-\mu,\sigma^2)$. Therefore, the overshoots and undershoots are symmetric, and we approximate them by the same expression. We use (31) and (37) to estimate

$$c\approx\frac{1}{\sigma\sqrt{2\pi}}\int_0^\infty e^{-(u-\mu)^2/(2\sigma^2)}\,du+\frac{-\mu}{2\left(\frac{\sigma^2+\mu^2}{-2\mu}-\sum_{n=1}^{\infty}\frac{1}{\sigma\sqrt{2\pi n}}\int_0^\infty ue^{-(u-n\mu)^2/(2n\sigma^2)}\,du\right)}. \tag{38}$$

Again, if our approximation is such that $c \geqslant 1$, we recommend using the backlog probability as a bound on the LCFS and loss probabilities. (We did not encounter such cases in our numerical experiments.)

Finally, using the technique in Siegmund (1985), the backlog probability when X has a normal$(-\mu, \sigma^2)$ distribution can be approximated for large d as

$$P(Y > d) \approx \exp\left(-\frac{2\mu}{\sigma^2}d\right)V(2\mu, \sigma) \approx \exp\left(-\frac{2\mu}{\sigma^2}d\right)\exp\left(2 * 0.583\mu/\sigma\right), \qquad (39)$$

where $V(\mu, \sigma) = 2\mu^{-2}\exp[-2\sum_{n=1}^{\infty} n^{-1}\Phi(-0.5|\mu/\sigma|n^{0.5})]$ ($\Phi(x)$ is the cumulative of a standard normal random variable) can be approximated as $\exp(-0.583|\mu|/\sigma)$ as $\mu/\sigma \to 0$.

5. Numerical Results

We designed the numerical experiments with three objectives in mind. The first objective is to show, based on simulation, that there might be substantial differences between the backlog, LCFS backlog, and loss probabilities. Such differences substantiate the importance of characterizing the LCFS backlog probability. The second objective is to investigate the accuracy of our analytic (up to the evaluation of the normal probability and its partial expectation in the normal case) approximations for the LCFS backlog probability. The third objective is to demonstrate the use of the LCFS backlog probability as a bound on the loss probability. To do this, we consider a production manager that chooses a base-stock level to minimize the holding cost subject to a constraint on the loss probability and show the potential savings from using the LCFS backlog probability in solving this problem.

We summarize the results of the numerical experiments as follows: The differences between the backlog, LCFS backlog, and the loss probabilities can be significant, e.g., the backlog probability can be more than twice the LCFS backlog probability, and the latter can be more than twice the loss probability. Our expressions for c were within -9% to $+5\%$ of the simulated c, and expressions for $P_{\text{LCFS}}(d)$ were a little less accurate when they used approximations for $P(Y > d)$. Furthermore, in line with the discussion at the end of §3, the accuracy of the approximation (31) improves with d, has no clear correlation with the drift, and improves with the standard deviation of X. Finally, using $P_{\text{LCFS}}(d)$ rather than $P(Y > d)$ as a bound on the loss probability led to an average of 31% cost savings.

We next discuss our experiments and their results in further detail.

5.1. The Experiments

We consider the three steps' distributions discussed in the last section. For each distribution, we consider negative drifts of 0.25, 0.5, and 0.75; and shortages probabilities, denoted by α, of 1%, 5%, 10%, and 20%. (In the inventory context, these probabilities correspond to service levels of 99%, 95%, 90%, and 80%, respectively.) For the normal distribution, we considered a standard deviation of $\sigma = 1$ and two scenarios with $\sigma = 2$. Thus, we consider a total of 60 scenarios: five steps' distributions (double- and shifted-exponential + three normals), times three drifts times four service levels. For each scenario, we consider seven different ds. (We present the inventory application and discuss how we chose these ds shortly.) Thus, we have a total of 420 instances for checking the differences between the backlog, LCFS backlog, and loss probabilities, and investigating the accuracy of our expressions for c and $P_{\text{LCFS}}(d)$.

For each instance, we simulated the random walk with $n = 500{,}000$ periods. In each simulation, we kept track of the backlog, LCFS backlog, and loss probabilities. Therefore, the 95% confidence intervals of the simulated probabilities are $\pm 1.96\sqrt{p(1-p)/n} < \pm 1.96\sqrt{1/n} < 0.28\%$ and are less than $1.5\% * p$ in most of our experiments.

For the inventory application, we let the holding cost be linear; thus, this cost is proportional to $d - E(\tilde{Y} \mid d)$, and the production manager chooses a base-stock level (or equivalently a threshold) d by solving

$$\min_{d} \quad (d - E(\tilde{Y} \mid d)) \qquad (40)$$

$$\text{s.t.} \quad P(\tilde{Y} = d) \leqslant \alpha. \qquad (41)$$

To solve (40), we considered two different methods. The first method is analytical. This method uses the backlog or the LCFS backlog probabilities as bounds on the loss probability; thus, using the expressions for $P(Y > d)$ and $P_{\text{LCFS}}(d)$ given in §4, it chooses base-stock levels

$$d\_BLG \equiv \min\{d \geqslant 0 \mid P(Y > d) \leqslant \alpha\} \quad \text{and}$$

$$d\_LCFS \equiv \min\{d \geqslant 0 \mid P_{\text{LCFS}}(d) \leqslant \alpha\},$$

respectively.

The second method is simulation based, and therefore it does not assure that the service-level constraint is met. We use the method suggested by Paschalidis and Liu (2003) in §5.2, and their discussion after the proof of Proposition 4. They recommend simulating the loss probabilities for a few ds, and then calculation of a convex piecewise-linear function mapping these ds to the simulated loss probability. Thus, given α, the optimal d is approximated using the line combining the thresholds of the two simulated loss probabilities around α. To compare our results with the ones in Paschalidis and Liu (2003), we calculated two base-stock levels, d. The first d is denoted by d_PL (for Paschalidis and Liu) and the second is denoted by d_B (for Baron).

Specifically, for each scenario, we simulated the loss probabilities for d_BLG, $0.75 * d\_BLG$, $0.5 * d\_BLG$, and then chose d_PL from the function mapping the loss probabilities to these ds. Similarly, we simulated the loss probabilities for d_LCFS and $d\_LCFS - (d\_BLG - d\_LCFS) = 2 * d\_LCFS - d\_BLG$ (if the latter is positive) and then chose d_B from the function mapping the loss probabilities to d_LCFS and $\max\{2*d\_LCFS-d\_BLG, 0\}$ (where the loss probability with $d=0$ is $P(X>0)$). This allowed 50% more simulation time to find d_PL than to find d_B.

Finally, to investigate the influence of the standard deviation, σ, on the accuracy of the approximations, we simulated the case of normal steps with $\sigma=2$ using the ds given by the case with $\sigma=1$. We do not consider the opposite case , i.e., using the optimal ds given by the case with $\sigma=2$ while letting $\sigma=1$ because it results in low backlog probabilities (smaller than 0.5% in most instances). Standard Monte Carlo simulation is not applicable in evaluating such rare events.

To summarize: For each scenario (i.e., steps' distribution, drift, and service level), we ran two sets of simulations. The first was with d_BLG, $0.75*d\_BLG$, $0.5*d\_BLG$, d_LCFS, and $2*d\_LCFS-d\_BLG$; and the second was with d_PL and d_B (where the last two values were calculated based on the first set of simulations, as explained above), for a total of seven different ds. Moreover, for each d, we also kept track of $E(\tilde{Y}\mid d)$ because it is required to evaluate the cost in (40).

5.2. Analysis of Numerical Results

Here we discuss the main findings from our numerical results. We then present the specific settings and results for each steps' distribution.

The simulations' results are reported in Tables 1–3 for double-exponential normal with $\sigma=1$, and normal with $\sigma=2$ steps, respectively. (The tables for normal steps with $\sigma=2$ and ds corresponding to $\sigma=1$ and for the shifted-exponential steps can be found in the electronic companion to this paper, available as part of the online version that can be found at http://or.pubs.informs.org/ecompanion.html.) For each scenario, we present d_BLG, d_LCFS, d_PL, and d_B along their corresponding simulated backlog, LCFS backlog, and loss probabilities; analytic backlog and LCFS backlog probabilities; and cost. (To save space, we do not present the corresponding information for $0.75*d\_BLG$, $0.5*d\_BLG$, and $2*d\_LCFS-d\_BLG$. Also, in the rest of the analysis we ignore 11 out of 420 instances where $d=0$.)

To measure the differences between the different probabilities, we calculated the ratios:

$$\text{backlog ratio} \equiv \frac{\text{simulated LCFS backlog probability}}{\text{simulated backlog probability}}$$

and

$$\text{loss ratio} \equiv \frac{\text{simulated loss probability}}{\text{simulated LCFS backlog probability}}.$$

We note that the backlog ratio is analagous to c in (18). To evaluate the accuracy of our expressions for c in (18), we define

$$c \text{ ratio} \equiv \frac{\text{average backlog ratio}}{\text{analytic } c},$$

where the average backlog ratio is calculated over the 28 ds with the same drift (seven ds for each of the four service levels). For each distribution and drift, in Table 4 we present the minimum, maximum, and average backlog ratio, the analytic c, and the c ratio. (The minimum and maximum backlog ratio were typically found for service level 99%, where the simulation is less accurate, or 80%, where the ds were small.) We observe that the backlog ratio is relatively fixed for the same drift independent of the ds. This supports the legitimacy of our assumption that $E(R)$ is independent of d. In contrast, the loss ratio varied more significantly for these cases. Furthermore, the average backlog and loss ratios were 0.60 and 0.70, respectively. Thus, the difference among the simulated backlog, LCFS backlog, and loss probabilities can be substantial for either of the distributions we considered. Moreover, the difference between the loss probabilities to the simulated LCFS backlog probabilities is less than the difference between the latter and the simulated backlog probabilities. Thus, using $P_{\text{LCFS}}(d)$ rather than $P(Y>d)$ as a bound on $P(\tilde{Y}=d)$ reduces the error by more than 50%. Moreover, as we might expect from Proposition 1, both ratios decrease with the drift.

It is seen that the c ratios in Table 4 are close to one in all cases, supporting the accuracy of our approximations. Moreover, other than for the double-exponential case (where the analytic c is the right one) and the normal with $\sigma=2$ using ds for normal with $\sigma=1$ (where the results are extremely accurate), the distance of the c ratio from one increases with the drift. However, because d is decreasing with the drift, we cannot attribute this improved accuracy solely to the increasing drift. Furthermore, comparing the c ratios of the three scenarios with normal steps, we see that the accuracy of the approximation for c is increasing with the standard deviation. This observation is in line with the discussion at the end of §3.

To evaluate the accuracy of our expressions for $P_{\text{LCFS}}(d)$, we define

$$\text{error } P_{\text{LCFS}}(d) \equiv 100*\frac{|\text{simulated } P_{\text{LCFS}}(d)-\text{approximated } P_{\text{LCFS}}(d)|}{\text{simulated } P_{\text{LCFS}}(d)}.$$

We observed that, on average, error $P_{\text{LCFS}}(d)$ is 5.13%. This accuracy includes many instances in which d is very small (e.g., for $\alpha=0.2$ and drift of 0.75); without these instances, the accuracy of our approximation is better. The accuracy of $P_{\text{LCFS}}(d)$ in the shifted-exponential case and for the two normal cases with $\sigma=2$ was as good as in the double-exponential case (when Proposition 1 gives the exact $P_{\text{LCFS}}(d)$); the accuracy was a little worse in the normal case with $\sigma=1$, partly due to the less accurate prediction of $P(Y>d)$.

For the solutions with d_BLG and d_LCFS, we also calculated the cost savings, defined as

$$\text{cost ratio} \equiv \frac{\text{cost with } d\_BLG-\text{cost with } d\_LCFS}{\text{cost with } d\_BLG}. \tag{42}$$

Table 1. Comparison of *d_BLG*, *d_LCFS*, *d_PL*, and *d_BA*; the simulated backlog, LCFS backlog, and loss probabilities; the analytic backlog and LCFS backlog probabilities; and the cost for random walks with double-exponential steps.

				Simulated probabilities of			Analytic probabilities of		
Service level (%)	Drift	*d*	Values	Backlog	LCFS backlog	Loss	Backlog	LCFS backlog	Cost
80	0.25	*d_BLG*	6.932	0.205	0.113	0.048	0.200	0.111	4.667
		d_LCFS	3.993	0.361	0.201	0.102	0.360	0.200	2.514
		d_PL	0.192	0.774	0.429	0.403	0.770	0.428	0.107
		d_B	2.270	0.513	0.285	0.172	0.508	0.282	1.363
	0.5	*d_BLG*	3.612	0.201	0.120	0.076	0.200	0.120	2.554
		d_LCFS	2.079	0.332	0.199	0.142	0.333	0.200	1.388
		d_PL	1.097	0.458	0.275	0.222	0.462	0.277	0.702
		d_B	1.502	0.403	0.242	0.184	0.404	0.242	0.979
	0.75	*d_BLG*	2.450	0.198	0.127	0.097	0.200	0.127	1.787
		d_LCFS	1.395	0.312	0.199	0.164	0.314	0.200	0.970
		d_PL	0.959	0.381	0.243	0.209	0.379	0.241	0.650
		d_B	1.105	0.355	0.226	0.191	0.356	0.226	0.757
90	0.25	*d_BLG*	10.397	0.099	0.055	0.022	0.100	0.056	7.505
		d_LCFS	7.458	0.181	0.101	0.042	0.180	0.100	5.082
		d_PL	3.238	0.418	0.232	0.126	0.419	0.233	2.001
		d_B	3.602	0.386	0.215	0.113	0.389	0.216	2.248
	0.5	*d_BLG*	5.691	0.101	0.061	0.036	0.100	0.060	4.276
		d_LCFS	4.159	0.165	0.099	0.062	0.167	0.100	2.993
		d_PL	2.991	0.247	0.148	0.099	0.246	0.148	2.065
		d_B	3.044	0.245	0.146	0.096	0.242	0.145	2.105
	0.75	*d_BLG*	4.067	0.100	0.064	0.046	0.100	0.064	3.143
		d_LCFS	3.012	0.157	0.100	0.074	0.157	0.100	2.250
		d_PL	2.462	0.199	0.127	0.096	0.199	0.127	1.798
		d_B	2.464	0.196	0.125	0.095	0.199	0.126	1.801
95	0.25	*d_BLG*	13.863	0.047	0.026	0.010	0.050	0.028	10.600
		d_LCFS	10.924	0.092	0.051	0.020	0.090	0.050	7.910
		d_PL	6.561	0.216	0.120	0.052	0.215	0.120	4.393
		d_B	5.669	0.260	0.144	0.065	0.257	0.143	3.717
	0.5	*d_BLG*	7.771	0.051	0.031	0.018	0.050	0.030	6.129
		d_LCFS	6.238	0.085	0.051	0.030	0.083	0.050	4.741
		d_PL	4.932	0.129	0.077	0.047	0.129	0.077	3.626
		d_B	4.771	0.135	0.081	0.050	0.136	0.082	3.500
	0.75	*d_BLG*	5.684	0.050	0.032	0.022	0.050	0.032	4.593
		d_LCFS	4.630	0.076	0.048	0.034	0.079	0.050	3.653
		d_PL	3.956	0.103	0.065	0.047	0.105	0.067	3.059
		d_B	3.905	0.109	0.070	0.050	0.107	0.068	3.005
99	0.25	*d_BLG*	21.910	0.010	0.006	0.002	0.010	0.006	18.022
		d_LCFS	18.971	0.018	0.010	0.004	0.018	0.010	15.361
		d_PL	14.947	0.038	0.021	0.008	0.040	0.022	11.533
		d_B	12.587	0.058	0.032	0.013	0.065	0.036	9.461
	0.5	*d_BLG*	12.599	0.010	0.006	0.003	0.010	0.006	10.714
		d_LCFS	11.067	0.016	0.009	0.005	0.017	0.010	9.217
		d_PL	9.449	0.028	0.017	0.010	0.029	0.017	7.692
		d_B	9.534	0.030	0.018	0.010	0.028	0.017	7.743
	0.75	*d_BLG*	9.440	0.010	0.007	0.004	0.010	0.006	8.174
		d_LCFS	8.385	0.016	0.010	0.007	0.016	0.010	7.144
		d_PL	7.670	0.022	0.014	0.009	0.021	0.014	6.468
		d_B	7.558	0.023	0.015	0.010	0.022	0.014	6.358

We calculated similar ratios for the solutions *d_B* and *d_PL* when both are feasible (with *d_PL* in the denominator). (The costs using *d_PL* and *d_B* are less relevant when these solutions are infeasible.)

We found that the average cost ratio is 31%. These savings decrease with the service level, but had no clear correlation with the drift. Also, the simulation-based solutions *dPL* and *dBA* were fairly good. Out of the total 48 cases

Table 2. Comparison of d_BLG, d_LCFS, d_PL, and d_BA; the simulated backlog, LCFS backlog, and loss probabilities; the analytic backlog and LCFS backlog probabilities; and the cost for random walks with normal steps where $\sigma = 1$.

				Simulated probabilities of			Analytic probabilities of		
Service level (%)	Drift	d	Values	Backlog	LCFS backlog	Loss	Backlog	LCFS backlog	Cost
80	0.25	d_BLG	2.636	0.202	0.121	0.072	0.200	0.119	1.828
		d_LCFS	1.595	0.334	0.201	0.137	0.337	0.200	1.050
		d_PL	0.912	0.469	0.281	0.220	0.474	0.281	0.577
		d_B	1.142	0.423	0.253	0.189	0.422	0.251	0.731
	0.5	d_BLG	1.026	0.199	0.137	0.118	0.200	0.131	0.775
		d_LCFS	0.606	0.293	0.199	0.180	0.305	0.200	0.443
		d_PL	0.496	0.321	0.218	0.200	0.340	0.223	0.359
		d_B	0.505	0.321	0.217	0.199	0.337	0.221	0.366
	0.75	d_BLG	0.490	0.179	0.135	0.129	0.200	0.140	0.398
		d_LCFS	0.252	0.239	0.178	0.173	0.286	0.200	0.199
		d_PL	0.103	0.280	0.206	0.204	0.358	0.250	0.080
		d_B	0.121	0.274	0.202	0.199	0.348	0.243	0.095
90	0.25	d_BLG	4.022	0.099	0.059	0.033	0.100	0.059	2.967
		d_LCFS	2.981	0.168	0.101	0.059	0.168	0.100	2.101
		d_PL	2.136	0.260	0.155	0.098	0.257	0.153	1.442
		d_B	2.163	0.254	0.152	0.096	0.253	0.151	1.462
	0.5	d_BLG	1.720	0.100	0.069	0.056	0.100	0.066	1.361
		d_LCFS	1.299	0.150	0.105	0.088	0.152	0.100	1.001
		d_PL	1.201	0.167	0.115	0.097	0.168	0.110	0.920
		d_B	1.195	0.169	0.117	0.099	0.169	0.111	0.914
	0.75	d_BLG	0.952	0.099	0.076	0.070	0.100	0.070	0.799
		d_LCFS	0.714	0.136	0.104	0.097	0.143	0.100	0.589
		d_PL	0.694	0.139	0.106	0.100	0.147	0.103	0.572
		d_B	0.694	0.139	0.106	0.100	0.147	0.103	0.572
95	0.25	d_BLG	5.409	0.050	0.030	0.016	0.050	0.030	4.190
		d_LCFS	4.368	0.083	0.050	0.028	0.084	0.050	3.264
		d_PL	3.404	0.134	0.080	0.046	0.136	0.081	2.446
		d_B	3.231	0.144	0.087	0.050	0.149	0.088	2.306
	0.5	d_BLG	2.413	0.051	0.035	0.028	0.050	0.033	1.984
		d_LCFS	1.992	0.076	0.053	0.042	0.076	0.050	1.602
		d_PL	1.850	0.089	0.061	0.049	0.088	0.058	1.475
		d_B	1.854	0.086	0.060	0.048	0.087	0.057	1.482
	0.75	d_BLG	1.414	0.050	0.039	0.035	0.050	0.035	1.225
		d_LCFS	1.176	0.072	0.056	0.051	0.071	0.050	1.003
		d_PL	1.207	0.069	0.053	0.049	0.068	0.048	1.031
		d_B	1.176	0.073	0.056	0.051	0.072	0.050	1.001
99	0.25	d_BLG	8.627	0.011	0.006	0.003	0.010	0.006	7.214
		d_LCFS	7.587	0.015	0.009	0.005	0.017	0.010	6.216
		d_PL	6.381	0.032	0.019	0.010	0.031	0.018	5.070
		d_B	6.070	0.036	0.021	0.011	0.036	0.021	4.799
	0.5	d_BLG	4.022	0.010	0.007	0.005	0.010	0.007	3.523
		d_LCFS	3.602	0.015	0.010	0.008	0.015	0.010	3.112
		d_PL	3.530	0.016	0.011	0.008	0.016	0.011	3.045
		d_B	3.430	0.018	0.013	0.010	0.018	0.012	2.945
	0.75	d_BLG	2.487	0.010	0.008	0.007	0.010	0.007	2.253
		d_LCFS	2.249	0.015	0.011	0.010	0.014	0.010	2.020
		d_PL	2.333	0.012	0.010	0.009	0.013	0.009	2.106
		d_B	2.249	0.015	0.012	0.010	0.014	0.010	2.021

(there is no service-level requirement for the case of normal steps with $\sigma = 2$ using the ds given by the case with $\sigma = 1$), the solutions based on dPL and dBA were feasible in 26 and 36 cases, respectively. For the 22 cases when both solutions were feasible, the cost of dBA was lower by 0.5% on average (measured similarly to the cost ratio). In general, dBA is a little more accurate, despite using less simulation time. This accuracy might be because the optimal d for the probability of loss was typically between d_LCFS and $\max\{2 * d\_LCFS - d\_BLG, 0\}$, whereas it

Table 3. Comparison of d_BLG, d_LCFS, d_PL, and d_BA; the simulated backlog, LCFS backlog, and loss probabilities; the analytic backlog and LCFS backlog probabilities; and the cost for random walks with normal steps where $\sigma = 2$.

Service level (%)	Drift	d	Values	Simulated probabilities of: Backlog	LCFS backlog	Loss	Analytic probabilities of: Backlog	LCFS backlog	Cost
80	0.25	d_BLG	11.710	0.204	0.112	0.04050	0.200	0.109	7.747
		d_LCFS	6.846	0.370	0.204	0.08910	0.367	0.200	4.215
		d_PL	0.000	0.834	0.450	0.45026	0.864	0.471	0.000
		d_B	3.782	0.542	0.298	0.16740	0.539	0.293	2.214
	0.5	d_BLG	5.272	0.197	0.118	0.07150	0.200	0.119	3.659
		d_LCFS	3.190	0.336	0.201	0.13740	0.337	0.200	2.101
		d_PL	1.942	0.456	0.272	0.21130	0.460	0.273	1.234
		d_B	2.282	0.421	0.251	0.18770	0.422	0.251	1.463
	0.75	d_BLG	3.126	0.200	0.130	0.09850	0.200	0.126	2.264
		d_LCFS	1.888	0.316	0.204	0.16850	0.318	0.200	1.311
		d_PL	1.475	0.364	0.234	0.20070	0.372	0.234	1.009
		d_B	1.536	0.356	0.229	0.19460	0.363	0.228	1.055
90	0.25	d_BLG	17.255	0.096	0.053	0.01760	0.100	0.054	12.236
		d_LCFS	12.391	0.183	0.101	0.03600	0.184	0.100	8.273
		d_PL	3.653	0.549	0.302	0.17140	0.547	0.298	2.137
		d_B	5.201	0.452	0.248	0.12180	0.451	0.246	3.118
	0.5	d_BLG	8.044	0.102	0.061	0.03340	0.100	0.059	5.915
		d_LCFS	5.963	0.166	0.100	0.05880	0.168	0.100	4.205
		d_PL	4.243	0.253	0.152	0.09650	0.259	0.154	2.875
		d_B	4.291	0.253	0.152	0.09640	0.256	0.152	2.908
	0.75	d_BLG	4.974	0.100	0.064	0.04520	0.100	0.063	3.807
		d_LCFS	3.737	0.162	0.104	0.07680	0.159	0.100	2.755
		d_PL	3.184	0.197	0.127	0.09630	0.196	0.123	2.309
		d_B	3.186	0.196	0.127	0.09650	0.196	0.123	2.313
95	0.25	d_BLG	22.800	0.052	0.028	0.00880	0.050	0.027	17.025
		d_LCFS	17.936	0.092	0.051	0.01650	0.092	0.050	12.781
		d_PL	9.075	0.270	0.148	0.05920	0.278	0.151	5.798
		d_B	7.145	0.341	0.188	0.08160	0.354	0.193	4.451
	0.5	d_BLG	10.817	0.051	0.030	0.01600	0.050	0.030	8.368
		d_LCFS	8.735	0.081	0.048	0.02670	0.084	0.050	6.538
		d_PL	6.836	0.135	0.081	0.04620	0.135	0.080	4.912
		d_B	6.689	0.141	0.084	0.04850	0.140	0.083	4.787
	0.75	d_BLG	6.823	0.050	0.032	0.02200	0.050	0.031	5.457
		d_LCFS	5.585	0.081	0.052	0.03630	0.080	0.050	4.335
		d_PL	4.868	0.104	0.067	0.04730	0.104	0.065	3.718
		d_B	4.860	0.106	0.068	0.04790	0.104	0.066	3.706
99	0.25	d_BLG	35.675	0.009	0.005	0.00150	0.010	0.005	29.010
		d_LCFS	30.811	0.020	0.011	0.00310	0.018	0.010	24.486
		d_PL	22.696	0.049	0.027	0.00830	0.051	0.028	16.935
		d_B	18.381	0.086	0.047	0.01520	0.087	0.047	13.176
	0.5	d_BLG	17.255	0.011	0.006	0.00340	0.010	0.006	14.420
		d_LCFS	15.173	0.018	0.011	0.00540	0.017	0.010	12.434
		d_PL	12.780	0.031	0.018	0.00960	0.031	0.018	10.184
		d_B	12.513	0.033	0.020	0.01080	0.033	0.019	9.934
	0.75	d_BLG	11.115	0.010	0.006	0.00420	0.010	0.006	9.539
		d_LCFS	9.877	0.017	0.011	0.00710	0.016	0.010	8.325
		d_PL	9.100	0.021	0.014	0.00910	0.021	0.013	7.586
		d_B	9.022	0.023	0.015	0.00990	0.022	0.014	7.502

was sometimes below $0.5 * d\_BLG$ and sometimes above $0.75 * d\_BLG$.

5.2.1. Double-Exponential Steps. We consider double-exponential steps with $\mu = 1$ and $1/\lambda$ of 1.25, 1.5, and 1.75. For any drift and service level, we found d_BLG using (33) and d_LCFS using (32) from Proposition 1. Because (32) is the exact expression for the LCFS backlog probability, the results in Table 1 show the accuracy of our simulation.

Table 4. Accuracy of the approximation for c based on (31).

Steps' distribution	Drift	Backlog ratio Min	Max	Average	Analytic c	c Ratio
Double-exponential	0.25	0.554	0.570	0.556	0.556	1.002
	0.5	0.594	0.607	0.600	0.600	1.001
	0.75	0.630	0.640	0.637	0.636	1.000
Shifted-exponential	0.25	0.451	0.466	0.456	0.461	0.990
	0.5	0.527	0.544	0.534	0.549	0.973
	0.75	0.592	0.615	0.605	0.631	0.958
Normal $\sigma = 1$	0.25	0.593	0.604	0.598	0.594	1.006
	0.5	0.665	0.696	0.688	0.657	1.047
	0.75	0.730	0.779	0.761	0.699	1.088
Normal $\sigma = 2$	0.25	0.546	0.556	0.550	0.549	1.001
	0.5	0.592	0.615	0.599	0.594	0.009
	0.75	0.631	0.649	0.645	0.629	1.026
Normal $\sigma = 2$ (with ds from normal $\sigma = 1$)	0.25	0.543	0.552	0.549	0.549	1.000
	0.5	0.581	0.600	0.594	0.594	0.999
	0.75	0.616	0.646	0.633	0.629	1.007

We observed that the average backlog and loss ratios are 0.6 and 0.62, respectively. Thus, the LCFS backlog probability might be substantially lower than the backlog probability, e.g., when the drift is 0.25 and $d\_BLG = 6.932$, the theoretical backlog probability is $P(Y > d) = 0.2$; however, the simulated LCFS backlog probability is 11.3% (about $P(Y > d)/2$) and the simulated loss probability is 4.8% (less than $P(Y > d)/4$). Moreover, the average c ratio is 1.001 and the average error $P_{\text{LCFS}}(d)$ is 1.9%; thus, the simulated LCFS backlog probabilities agree with the analytic ones in (32). Finally, the cost ratio ranged between 12.6% and 46.1%, averaged at 28.2%.

5.2.2. Shifted-Exponential Steps. We consider shifted-exponential steps with $\mu = 1$ and a shift r of 1.25, 1.5, and 1.75. For any drift and service level, we found d_BLG from (35) and d_LCFS from (36). The detailed results are reported in Table 1 of the electronic companion to this paper.

We found that the average backlog and loss ratios are 0.53 and 0.75, respectively, and the average c ratio is 0.974, which is translated to an average error $P_{\text{LCFS}}(d)$ of 3.1% because $P(Y > d)$ is exact. The cost ratio ranged between 14.5% and 100% (100% for service level 80% and a drift of 0.75), averaged at 43.7%.

5.2.3. Normal Steps. We consider normal steps with μ of −0.25, −0.5, and −0.75 with a standard deviation of $\sigma = 1$ or $\sigma = 2$. For any μ and service level, we found d_BLG from (39) and d_LCFS by multiplying c from (38) by $P(Y > d)$ from (39). We only evaluated the first 250 elements of the infinite series in (38) to express c.

For a concrete example of the inventory problem, one can think of a normal$(10, \sigma)$ demand process (thus, the probability of negative demand is negligible) and deterministic capacity $r \in (10.25, 10.5, 10.75)$. The random walk associated with these demands and capacities are identical to the ones we consider.

For the $\sigma = 1$ case, the results are reported in Table 2. We found that the average backlog and loss ratios are 0.68 and 0.79, respectively; the average c ratio is 1.047; and the average error $P_{\text{LCFS}}(d)$ is 5.4%. Note that for $\mu = -0.75$ and service level of 80%, the resulting ds are relatively small; thus, the approximation for $P(Y > d)$ and for c are poor. Finally, the cost ratio ranged between 10.3% and 49.8%, averaged at 26%.

For the $\sigma = 2$ case, the results are reported in Table 3. We found that the average backlog and loss ratios are 0.6 and 0.58, respectively, the c ratio is 1.012, and the average error $P_{\text{LCFS}}(d)$ is 2.2%. Finally, the cost ratio ranged between 12.7% and 45.6%, averaged at 27%.

To compare the change in accuracy of our approximations with the standard deviation of X, we ran the case with normal steps with $\sigma = 2$ using the optimal ds for normal steps with $\sigma = 1$. The detailed results are reported in Table 2 of the electronic companion to this paper. We saw that the average backlog and loss ratios are 0.59 and 0.78, respectively, the average c ratio is 1.002, and the average error $P_{\text{LCFS}}(d)$ is 13%. Thus, in this scenario, with relatively small ds and high backlog probabilities, our approximations for c are better than the approximation for $P(Y > d)$.

6. Summary

This paper uses the sample path approach to relate the backlog, LCFS backlog, and loss probabilities defined for inventory systems represented using regulated random walks. These probabilities are plausible service-level measures for systems that can be modeled by one- or two-sided regulated random walks such as single-server queues, and token-bucket admission controls. We expressed the LCFS backlog probability as the backlog probability times a coefficient c that is smaller than one, and showed that the LCFS backlog probability is an upper bound on the loss probability.

Analytic expressions for the LCFS backlog probability combine known expressions for the backlog probability with expressions developed for c. We express c in closed form for random walks with double-exponential steps and suggest approximating it using (31) when the steps are from a nonarithmetic distribution, and have a conjugate point.

A comparison of our analytic results with simulation results for random walks with double-exponential, shifted-exponential, and normal steps shows that our approximations of the LCFS backlog probability are on average 5.13% of this simulated probability. Moreover, the distance of the simulated loss probabilities from the LCFS backlog probabilities is less than the distance between the latter and the backlog probability. Thus, using the LCFS backlog probability rather than the backlog probability as a bound on the loss probability reduced the error by more than 50%.

Our methods could result in costs savings when planning resources for systems whose workload can be modeled using a one- or two-sided regulated random walk. We demonstrate this in a simple inventory model, where our results reduce inventory levels and associated costs while maintaining a required service level. In these settings, our approximations led to an average cost savings of 31%. These results could also support managerial decisions on a backlog allocation rule (i.e., allocating backlog according to FCFS or to LCFS).

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.pubs.informs.org/ecompanion.html.

Acknowledgments

This paper is based to a large extent on the author's Ph.D. thesis. The author is thankful to his advisor Gabriel R. Bitran and Ph.D. committee member Dirk Beyer for very helpful discussions and comments on earlier versions of this paper. The author thanks Yashan Wang for his help with the proof of Lemma 1, and David Perry and Joseph Milner for helpful comments on an earlier version of this paper. The author also thanks the two anonymous reviewers and the associate editor for remarks that substantially improved this paper. This research was partially supported by a grant from the Natural Sciences and Engineering Research Council of Canada.

References

Abate, J., L. G. Choudhury, W. Whitt. 1995. Exponential approximation for tail probabilities in queues, I: Waiting times. *Oper. Res.* **43**(5) 885–901.

Asmussen, S. 2003. *Applied Probability and Queues*, 2nd ed. Springer, New York.

Baron, O. 2003. Pricing and admission control for shared computer services using the token bucket mechanism. Ph.D. thesis, Sloan School of Management, Massachusetts Institute of Technology, Boston, MA.

Baron, O., D. Beyer, G. R. Bitran. 2005. Pricing of shared computer services. *J. Revenue Pricing Management* **4**(1) 54–65.

Berger, W. A. 1991. Performance analysis of a rate-control throttle when tokens and jobs queue. *IEEE J. Selected Areas Comm.* **9**(2) 165–170.

Berger, W. A., W. Whitt. 1992. The Brownian approximation for rate-control throttle and the G/G/1/C queue. *Discrete Event Dynam. System: Theory Appl.* **2** 7–60.

Bradley, J. R., P. W. Glynn. 2002. Managing capacity and inventory jointly in manufacturing systems. *Management Sci.* **48**(2) 273–288.

Cohen, J. W. 1982. *The Single Server Queue*, 2nd ed. North-Holland, Amsterdam, The Netherlands.

de Kok, A. G. 1985. Approximations for lost-sales production/inventory control model with service level constraints. *Management Sci.* **31**(6) 729–737.

Gallager, G. R. 1996. *Discrete Stochastic Process*. Kluwer Academic Publishers, Boston, MA.

Glasserman, P. 1997. Bounds and asymptotics for planning critical safety stock. *Oper. Res.* **45**(2) 244–257.

Glasserman, P., S. Tayur. 1996. A simple approximation for a multi-stage capacitated production-inventory system. *Naval Res. Logist.* **43** 41–58.

Harrison, J. M. 1985. *Brownian Motion and Stochastic Flow Systems*. Wiley, New York.

Kushner, J. H., P. Dupuis. 2001. *Numerical Methods for Stochastic Control Problems in Continous Time*. Springer, New York.

Paschalidis, C. I., Y. Liu. 2003. Large deviations-based asymptotics for inventory control in supply chains. *Oper. Res.* **51**(3) 437–460.

Prabhu, N. U. 1980. *Stochastic Storage Processes Queues, Insurance Risk and Dams*. Springer-Verlag, New York.

Ross, M. S. 1974. Bounds on the delay distribution in GI/G/1 queues. *J. Appl. Probab.* **11** 417–421.

Ross, M. S. 1983. *Stochastic Processes*. Wiley Series in Probability and Mathematical Statistics, Wiley, New York.

Siegmund, D. 1985. *Sequential Analysis: Tests and Confidence Intervals*. Springer, New York.

Whitt, W. 2002. *Stochastic-Process Limits*. Springer, New York.

Woodroofe, M. 1979. Repeated likelihood ratio tests. *Biometrica* **16** 453–463.