Approximating the Distribution of a Sum of Log-normal Random Variables

Barry R. Cobb
Virginia Military Institute
Lexington, VA, USA
cobbbr@vmi.edu

Rafael Rumí
Antonio Salmerón
Universidad de Almería, Spain
rrumi,antonio.salmeron@ual.es

Abstract

This paper introduces a process for estimating the distribution of a sum of independent and identically distributed log-normal random variables (RVs). The procedure involves using the Fenton-Wilkinson method to estimate the parameters for a single log-normal distribution that approximates the sum of log-normal RVs. Once these parameters are determined, a mixture of truncated exponentials (MTE) function is determined to approximate this distribution. The MTE parameters are stated as polynomial functions of the log-normal scale parameter. Applications to inventory management are presented that demonstrate the usefulness of the MTE approximation.

1 Introduction

Finding the probability density function (PDF) for a sum of log-normally distributed random variables (RVs) is an important problem in business and telecommunications (Beaulieu et al., 1995). This paper proposes a tractable approximation to the PDF for a sum of log-normal RVs that can be utilized in Bayesian networks (BNs) and influence diagrams (IDs).

Consider the following business application in inventory management. Suppose the independent and identically distributed (i.i.d.) values of customer demand for a business, X_ℓ, in each period $\ell = 1, \ldots, L$, are log-normally distributed, i.e. $X_\ell \sim LN(\mu_{X_\ell}, \sigma^2_{X_\ell})$. This means that total customer demand over the L periods, X, is determined as the following sum of i.i.d. RVs:

$$X = X_1 + X_2 + X_3 + \ldots + X_L \ . \quad (1)$$

The value L represents the fixed lead time between the placement and arrival of an inventory order. The business needs to establish order quantity and reorder point policies that minimize inventory costs. Finding a tractable approximation to the distribution for X is important to solving this problem.

Fenton (1960) and Schwartz and Yeh (1982) estimate the PDF for a sum of log-normal RVs using another log-normal PDF with the same mean and variance. The Fenton approximation (sometimes referred to as the *Fenton-Wilkinson* (FW) method) is simpler to apply, and for a wide range of log-normal parameters has been shown to be reasonably accurate in comparison to the Schwartz-Yeh (SY) method (Beaulieu et al., 1995).

While the FW method produces a PDF that in some cases is a good approximation to the PDF for a sum of log-normals, this still does not result in a tractable representation that can be incorporated in BNs and IDs. This is because such models require a functional form that can be combined with other functions, then integrated in closed-form. This paper introduces a new mixtures of truncated exponentials (MTE) approximation to the log-normal PDF. MTE functions were introduced by Moral et al. (2001) to facilitate closed-form mathematical calculations in BNs.

One technique (Cobb et al., 2006) suggested for approximating a log-normal PDF with an MTE function requires a nonlinear optimization problem to be solved for each possible combination of μ and σ^2. In this paper, an approximation is introduced that calculates the MTE parameters as polynomial functions of only the scale parameter, σ, for a range of values needed to model the inventory application. This reduces the number of tabular inputs required to efficiently incorporate the FW method into an inventory management system using BNs or IDs.

The next section gives definitions and notation. Section 3 describes the MTE approximation. Section 4 applies the MTE approximation to inventory management examples and compares the results to those obtained with other techniques. The final section summarizes and gives suggestions for future research.

2 Definitions

Variables in this paper will be denoted by capital letters, e.g., X, Y, Z, with specific values shown in lower-case, e.g., x, y, z. The state space of X is denoted by Ω_X. The expressions $\exp(x)$ and e^x are used interchangeably.

2.1 Log-normal Distribution

A RV X is log-normal, i.e. $X \sim LN(\mu, \sigma^2)$, if and only if $\ln(X) \sim N(\mu, \sigma^2)$. A log-normal RV has the PDF

$$f_X(x) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp\left(\frac{-(\ln x - \mu)^2}{2\sigma^2}\right) \ , \ x > 0 \ ,$$

for any $\sigma^2 > 0$. The expected value of X is $E(X) = \exp(\mu + 0.5\sigma^2)$ and the variance of X is $Var(X) = \big(\exp(\sigma^2) - 1\big)\exp(2\mu + \sigma^2)$.

The mode of $f_X(x)$ is $m = \exp(\mu - \sigma^2)$ and the inflection points to the right and left of the mode are $d^{\pm} = \exp\left(\frac{1}{2}(2\mu - 3\sigma^2 \pm \sigma\sqrt{4+\sigma^2})\right)$.

If $X \sim LN(\mu, \sigma^2)$, then $aX \sim LN(\mu + \ln(a), \sigma^2)$ where $a > 0$. Conveniently, then, we can find a PDF for a $X \sim LN(\mu, \sigma^2)$ as a convolution of $Y \sim LN(0, \sigma^2)$ as follows:

$$f_X(x) = (1/\exp(\mu)) \cdot f_Y\left((1/\exp(\mu)) \cdot x\right) \quad (2)$$

2.2 Mixtures of Truncated Exponentials

The functional form of the log-normal PDF does not permit integration in closed-form. This makes this PDF difficult to incorporate into BNs and IDs. One approach that can allow log-normal RVs to be included in these models is to approximate log-normal PDFs in the BN or ID with MTE potentials.

MTE Potential (Moral et al., 2001). Let X be a continuous chance variable. Given a partition $\Omega_1, \ldots, \Omega_n$ that divides Ω_X into hypercubes, an n-piece, m-term MTE potential $\phi : \Omega_X \mapsto \mathcal{R}^+$ has components

$$\phi_h(x) = a_{1h} + \sum_{i=1}^{m} a_{2i,h} \exp(a_{2i+1,h} \cdot x)$$

for $h = 1, \ldots, n$, where $a_{jh}, j = 1, \ldots, 2m+1$ are real numbers. Thus, an n-piece, m-term MTE potential requires $2mn + n$ parameters.

MTE potentials can approximate both PDFs and utility functions. The optimization procedures outlined by Cobb et al. (2006) and Langseth et al. (2012) can be used to determine the parameters (the values a_{jh}) required to approximate PDFs with MTEs.

2.3 Fenton-Wilkinson (FW) Approximation

Consider the sum of L i.i.d. log-normal RVs, X, as specified in (1) where each $X_\ell \sim LN\left(\mu_{X_\ell}, \sigma^2_{X_\ell}\right)$ with the expected value and variance described in Section 2.1. The expected value and variance of X are $E(X) = L \cdot E(X_\ell)$ and $Var(X) = L \cdot Var(X_\ell)$. The FW approximation is a log-normal PDF with parameters μ_X and σ^2_X such that $\exp(\mu_X + 0.5\sigma^2_X) = L \cdot E(X_\ell)$ and

$$\big(\exp(\sigma^2_X) - 1\big) \cdot \exp(2\mu_X + \sigma^2_X) = L \cdot Var(X_\ell) \ .$$

Solving for μ_X and σ^2_X gives

$$\sigma^2_X = \ln\big((\exp(\sigma^2_{X_\ell}) - 1)/L + 1\big) \quad (3)$$

and

$$\mu_X = \ln(L \cdot \exp(\mu_{X_\ell})) + 0.5(\sigma^2_{X_\ell} - \sigma^2_X) \ . \quad (4)$$

Example 1. Suppose X_ℓ, $\ell = 1, \ldots, 5$ are distributed as $X_\ell \sim LN(0.69, 1.07^2)$. The sum, X, of these i.i.d. RVs has expected value $E(X) = 5 \cdot \exp(0.694 + 0.5 \cdot 1.074^2) = 17.8$ and $Var(X) = 5 \cdot Var(X_\ell) = 137.5$. Applying the FW approximation entails using the $LN(2.7, 0.6^2)$ distribution as an estimated PDF for X. Note that this distribution has an expected value of $\exp(2.7 + 0.5 \cdot 0.6^2) = 17.8$ and a variance of $(e^{0.6^2} - 1)(e^{2\cdot 2.7}e^{0.6^2}) = 137.5$.

3 MTE Approximation to the Log-normal Distribution

This section describes some enhancements to the MTE approximation to the log-normal PDF detailed by Cobb et al. (2006). Specifically, the previous model required the MTE parameters a_{jh} to be determined for each combination of μ and σ^2. However, by applying the convolution in (2), we only need to estimate an MTE density function for each value of σ.

Note from (3) that σ_X^2 will be smaller than $\sigma_{X_\ell}^2$ and will decrease as L increases. For instance, for $\sigma_{X_\ell}^2 = 2$, $\sigma_X^2 < 1$ if $L \geq 4$. The FW approximation has been shown to be adequate when $\sigma_{X_\ell}^2 < 2$ (Beaulieu et al., 1995). Additionally, for inventory management problems, such as the one mentioned earlier, Das (1983) has observed that $\sigma_X^2 < 1$ normally holds. Thus, focusing on a range for σ_X in the interval $[0.05, 1]$ will be adequate for the problems in Section 4.

3.1 Regions for σ and Domain Splits

Four *regions* for the log-normal parameter σ are established as: $R_1 : [0.05, 0.1)$, $R_2 : [0.1, 0.2)$, $R_3 : [0.2, 0.63)$, and $R_4 : [0.63, 1]$. The values $\sigma = 0.1$ and $\sigma = 0.2$ are chosen because the shape of the log-normal PDF changes rapidly for σ values near zero. The value $\sigma = 0.63$ is chosen as the value to divide R_3 and R_4 because it is the point where $d^- \approx \exp(-2\sigma)$, thus this value determines the domains of two pieces of the MTE potential, as shown below.

Suppose we have a RV $Y \sim LN(0, \sigma^2)$. We will use a 5-piece, 2-term MTE approximation to the PDF for Y. The domains of the five pieces are determined using:

$$\begin{array}{ll} y_0 : \exp(-3\sigma) & y_1 : \mathrm{Max}\{d^-, \exp(-2\sigma)\} \\ y_2 : m = \exp(-\sigma^2) & y_3 : d^+ \\ y_4 : \exp(\sigma) & y_5 : \exp(3\sigma) \end{array}$$

3.2 MTE Specification

The un-normalized 5-piece, 2-term MTE approximation to the $LN(0, \sigma^2)$ PDF is

$$\hat{f}_Y^{(u)}(y) = \begin{cases} \hat{a}_{1h}(\sigma) + \\ \sum_{i=1}^{2} \hat{a}_{2i,h}(\sigma) \exp(\hat{a}_{2i+1,h}(\sigma)(y - m(\sigma))) \\ \qquad \text{if } y_{h-1} \leq y < y_h \end{cases} \tag{5}$$

for $h = 1, \ldots, 5$. All functions are assumed to equal zero in undefined regions. By adding pieces, or by fitting more terms in the first and fifth pieces, the domain of the MTE function could be extended, at the expense of additional computational burden in making calculations with the resulting distribution (Cobb et al., 2006; Romero et al., 2006). Any selection of number of pieces and terms in an MTE function involves some trade-off between accuracy and computational burden.

3.3 Functional Representation of MTE Parameters

To reduce the tabular inputs required to completely define the family of MTE approximations to the $LN(0, \sigma^2)$ PDFs, we use approximate parameters $\hat{a}_{jh}$. These parameters are determined using the polynomial functions

$$\hat{a}_{jh}(\sigma) = \sum_{k=1}^{5} \hat{c}_{kjh}^{(r)} \sigma^{k-1} \tag{6}$$

for $h = 1, \ldots, 5$ and $r = 1, \ldots, 4$. The value of r is determined according to the region – R_1, R_2, R_3, or R_4 – where σ is located. Thus, to define the parameters for a 5-piece, 2-term MTE distribution where σ is located in any of these four regions requires a table of $4 \cdot 5 \cdot (2 \cdot 2 \cdot 5 + 5)$ $= 500$ constants, $\hat{c}_{kjh}^{(r)}$.

To estimate the constants $\hat{c}_{kjh}^{(r)}$, we use the Cobb et al. (2006) nonlinear optimization procedure to find the "exact" MTE parameters a_{jh} for $0.05 \leq \sigma \leq 1$ in increments of 0.01 for σ. Linear regression is then used to estimate the model in (6). As an example, Figure 1 shows a plot of the actual MTE parameter a_{24} for values $0.63 \leq \sigma \leq 1$, as determined using the nonlinear optimization procedure. The equation

$$\begin{aligned} \hat{a}_{24}(\sigma) = {} & 101.8201 - 2.5357\sigma \\ & + 3.2139\sigma^2 - 1.9887\sigma^3 + 0.5033\sigma^4 \end{aligned}$$

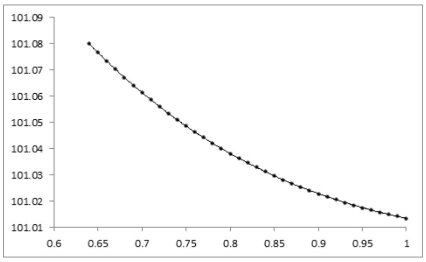

Figure 1: Plots of a_{24} (dots) and $\hat{a}_{24}$ (line).

is estimated using least squares regression and can be used to find one of the MTE parameters required to approximate any log-normal PDF with $\sigma \in R_4$. This equation fits the actual data ($R^2 = 1$ to four decimal places) very closely and is represented by the curve connecting the dots in Figure 1.

Once the constants $\hat{c}^{(r)}_{kjh}$ are stored, no further use of the nonlinear optimization procedure is required and a log-normal PDF with any μ value and any $\sigma \in [0.05, 1]$ can be approximated.

3.4 Summarized MTE Approximation Process

Ultimately, we want to approximate the distribution for $X = X_1 + \cdots X_L$ where X_ℓ, $\ell = 1, \ldots, L$ are i.i.d. $LN(\mu_{X_\ell}, \sigma^2_{X_\ell})$ RVs. Given the formulation we have described for approximating the $LN(0, \sigma^2)$ PDF, the process can be summarized as follows:

1. Calculate the FW parameters σ^2_X and μ_X from expressions (3) and (4).
2. Calculate the MTE parameters $\hat{a}_{jh}(\sigma_X)$ using equation (6).
3. Construct an un-normalized version, $\hat{f}^{(u)}_Y(y)$, of the MTE approximation to the PDF for a RV Y that has the $LN(0, \sigma^2_X)$ distribution using (5). Although the "exact" parameters $a_{jh}(\sigma)$ were determined for a PDF that integrated to 1, because the $\hat{a}_{jh}(\sigma)$ are approximate parameters, the result of expression (5) is not guaranteed to integrate to 1.
4. Normalize the PDF for Y by calculating $k_Y = \int_{-\infty}^{\infty} \hat{f}^{(u)}_Y(y)\, dy$ and finding $\hat{f}_Y(y) = k_Y^{-1} \cdot \hat{f}^{(u)}_Y(y)$.
5. Use the convolution operation in (2) to create the normalized PDF $\hat{f}_X(x)$ for X.

Example 2. Continuing from Example 1, to construct an MTE approximation to the $LN(2.7, 0.6^2)$ PDF for X, we first need to build the MTE approximation to the $LN(0, 0.6^2)$ PDF. Table 1 shows the constants $c^{(3)}_{kj1}$ needed to determine the parameters for the first (of five) pieces of this MTE potential. The second row shows σ^{k-1}, $k = 1, \ldots, 5$, while the other rows show the constants need to calculate each of the five parameters $\hat{a}_{j1}$. These parameters are calculated in the last column as the sumproduct of the σ^{k-1} row and the corresponding row of $\hat{c}^{(3)}_{kj1}$'s.

The un-normalized MTE approximation to the $LN(0, 0.6^2)$ (with mode $m = 0.698$) PDF is

$$\hat{f}^{(u)}_Y(y) = \begin{cases} -0.098 + 87.823e^{10.208(y-m)} & \\ -113.664e^{11.609(y-m)} & 0.165 \le y \le 0.311 \\ -8.727 - 100.994e^{-0.760(y-m)} & \\ +110.515e^{-0.694(y-m)} & 0.311 \le y \le 0.698 \\ \vdots & \\ 0.001 - 93.350e^{-1.749(y-m)} & \\ +94.235e^{-1.744(y-m)} & 1.822 \le y \le 6.050\ . \end{cases}$$

The normalization constant k_Y is 0.971844, so the normalized MTE distribution is $\hat{f}_Y(y) = \hat{f}^{(u)}_Y(y)/0.971844$. The MTE approximation to the $LN(2.7, 0.6^2)$ PDF for X is

$$\hat{f}_X(x) = (1/14.8797) \cdot \hat{f}_Y\left((1/14.8797) \cdot x\right)\ .$$

The result is shown in Figure 2 overlaid on the actual $LN(2.7, 0.6^2)$ PDF along with two other distributions used in Section 4 (it is the distribution with the second highest mode).

4 Inventory Management Examples

In this section we consider the problem of establishing a (Q, R) policy in a continuous review inventory system. In this model, Q is the order quantity and R is the inventory level when an order is placed, or *reorder point*. The expected cost function (Hadley and Whitin, 1963) is

Table 1: Calculation of parameters for an MTE distribution approximating the $LN(0, 0.6^2)$ PDF.

k	5	4	3	2	1	
σ^{k-1}	0.1296	0.216	0.36	0.6	1.0000	$\hat{a}_{j1}$
$\hat{c}^{(3)}_{k11}$	-12.8729	19.2704	-10.3045	2.2788	-0.2493	-0.098
$\hat{c}^{(3)}_{k21}$	218.6796	-378.9597	241.8550	-68.0000	95.0700	87.823
$\hat{c}^{(3)}_{k31}$	606.9384	-1120.0537	785.4244	-252.0248	41.9427	10.208
$\hat{c}^{(3)}_{k41}$	185.4334	-322.1284	206.3402	-58.3289	-107.4016	-113.664
$\hat{c}^{(3)}_{k51}$	852.5394	-1556.6422	1075.2962	-338.6557	53.4414	11.609

$$TC(q,r) = \frac{KY}{q} + \frac{\pi Y S_R(r)}{q} + c\left(\frac{q}{2} + r - E(X)\right) . \tag{7}$$

In this formula, K is the fixed cost per order, Y is the annual demand, c is the unit holding cost per year, and π is the stockout cost per unit. The terms are the annual ordering, stockout, and holding costs, respectively. The third term assumes the firm holds an average *safety stock* of $R - E(X)$. The shape of the distribution for X enters the model via the expected shortage per cycle $S_R(r)$. This is approximated as

$$\hat{S}_R(r) = \int_r^\infty (x - r) \cdot \hat{f}_X(x)\, dx \tag{8}$$

where $\hat{f}_X(x)$ is an approximation to the lead time demand (LTD) distribution.

4.1 Uncertain Demand

Das (1983) presents an approximate method for finding a (Q, R) policy where LTD is log-normal. Presumably, an approximately log-normal PDF for LTD could be the result of demand per unit time for a fixed lead time being log-normal so that LTD is a sum of log-normals that can be approximated using the FW method. In other words, Das presents results that can be used to find an approximately optimal Q and R after σ_X and μ_X from (3) and (4) have been determined. The test cases in Table 2 are used by Das, with $K = 30$ in each case, $\pi = 5$ for Cases 1 and 2, and $\pi = 6$ for Case 3.

Das specifies the values in the μ_X and σ_X columns of Table 2. The values in the μ_{X_ℓ} and σ_{X_ℓ} columns are calculated assuming that lead time is fixed at $L = 5$ and that LTD is the i.i.d. sum of $L = 5$ one period demand values that are log-normally distributed as $LN(\mu_{X_\ell}, \sigma^2_{X_\ell})$. The

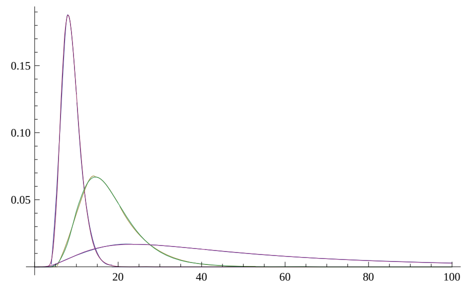

Figure 2: MTE approximations and actual log-normal PDFs for the Das (1983) test cases.

Q_D and R_D columns contain the approximately optimal solutions obtained by Das.

To use the MTE approximation to find optimal Q and R for the three test cases, we approximate the three log-normal PDFs with μ_X and σ_X shown above using the process from Section 3.4. These three MTE distributions are shown in Figure 2, overlaid on the actual log-normal PDFs. The function with the largest mode corresponds to $\mu_X = 1.6$ and $\sigma_X = 0.8$, whereas the function skewed the farthest right approximates the $LN(2.3, 1)$ PDF.

Using the MTE approximations as the $\hat{f}_X(x)$ PDFs for LTD allows us to establish a closed-form expression for $\hat{S}_R(r)$ in (8). For example, in Case 1, the expression is

$$\hat{S}_R(r) = \begin{cases} 78381.9 - 429.776r + 0.809782r^2 \\ +75472.7e^{-0.0101464r} - 153835e^{-0.00776065r} \\ \qquad 17.2631 \le r < 27.1126 \\ 0.228577 - 0.00531733r + 0.0000306374r^2 \\ -1582.93e^{-0.117546r} + 1601.78e^{-0.117187r} \\ \qquad 27.1126 \le r \le 90.0171\ . \end{cases} \tag{9}$$

Table 2: Parameters and solutions for the Das (1983) test cases.

Case	Y	μ_X	σ_X	μ_{X_ℓ}	σ_{X_ℓ}	h	(Q^*, R^*)	(Q_D, R_D)	(Q_M, R_M)	HD_D	HD_M	CT_M
1	400	2.7	0.6	0.69	1.07	4	(90.0,24.8)	(89.7,25.8)	(88.6,25.2)	0.03%	0.05%	1.4s
2	100	1.6	0.8	−0.54	1.30	2	(62.0,8.1)	(61.9,9.5)	(60.7,8.6)	0.24%	0.00%	1.5s
3	300	2.3	1	0.06	1.50	3	(105.4,23.7)	(104.7,25.5)	(95.8,26.0)	0.05%	0.18%	1.3s

Once the $\hat{S}_R(r)$ expression is established we use it to approximate $TC(q,r)$ in (7), and we can subsequently find the values for Q and R that minimize $\hat{TC}(q,r)$. We perform this optimization in Mathematica 8.0. All computations were done on a PC with 16GB of memory and an Intel Core2 Duo 2.93 GHz processor.

For comparison purposes, a simulation-optimization technique is employed to find the "exact" solution. Using Oracle Crystal Ball software's tabu-search driven OptQuest tool, this approach tests numerous possible values for Q and R. Demand is simulated for each day, then total cost is calculated by computing shortage cost as $\text{Max}\{0, R - X\} \cdot (\pi Y/Q)$ on each simulation trial. To find an accurate solution, we used 10,000 trials to test each selected (Q, R) combination and ran OptQuest for several hours. For Case 1, the exact solutions are determined to be $Q^* = 90$, $R^* = 24.8$. The expected total cost using a longer simulation of 1,000,000 trials with the optimal values is 390.7.

To judge the effectiveness of the MTE approach versus the exact solutions, the cost differential used by Heuts et al. (1986) can be calculated as

$$HD = 100\% \cdot \frac{TC(Q_M, R_M) - TC(Q^*, R^*)}{TC(Q^*, R^*)} . \tag{10}$$

This expression uses the MTE solutions, Q_M and R_M, and simulates the expected total cost calculated over 1,000,000 trials. Thus, HD measures the error that occurs when the MTE solutions are employed as an approximation to the true solutions. The HD cost differential is calculated similarly for Q_D and R_D.

Solutions found in all three test cases defined by Das are shown in Table 2 (M and D subscripts correspond to the MTE and Das methods, respectively). Computing time for the MTE method, CT_M, is also reported. Das does not report computing time required to obtain solutions using the analytical method, but does state that they are about 20% of that required for an iterative method due to Hadley and Whitin (1963). Actual computational times for the Das method are likely negligible.

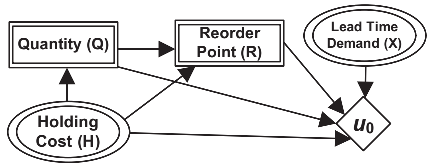

Figure 3: Influence diagram model.

The Das solution has a lower HD in two cases, whereas the MTE solution is more accurate in the second case. In summary, the MTE method performs comparably with the Das method and the simulation-optimization method for these problems and uses minimal computation time.

4.2 Uncertain Demand and Holding Cost

Another motivation for approximating the LTD distribution with an MTE potential is to use IDs designed for decision-making with continuous variables. In such models, the functional form of the PDFs and utility functions must be such that the mathematical operations of addition, multiplication, and integration maintain a closed-form result that can be used in further calculations.

Suppose that we have a problem similar to the Das Case 1 example, but that holding cost is a RV because the company finances inventory purchases with a variable interest rate line of credit. Before each order, the company knows the true current holding cost and can adjust its order accordingly. The ID model for this problem appears in Figure 3. Chance, decision, and utility nodes appears as ovals, rectangles, and diamonds, respectively.

The utility function in the problem is

$$u_0(q,r,h) = -\frac{K}{Y} \cdot g_1(q) - \pi \cdot Y \cdot \hat{S}_R(r) \cdot g_1(q) \\ -g_2(h) \cdot (0.5 \cdot g_3(q) + g_4(r) - E(X)) \ , \tag{11}$$

where $g_1(q) = 1/q$, $g_2(h) = h$, $g_3(q) = q$, and $g_4(r) = r$. This is similar to $TC(q,r)$ in (7), except that negative signs have been placed in front of each term and holding cost is now a RV. The utility function is not explicitly a function of X because the term $S_R(r)$ is created by integrating out X. We include X in the graphical ID to emphasize that the PDF for X affects the total cost through $\hat{S}_R(r)$. The computation that creates the function $\hat{S}_R(r)$ becomes a marginalization operation for X.

The LTD distribution is the same MTE distribution used in Section 4.1 for Case 1. Holding cost, H, is $N(4, 4/9)$, and is approximated with an MTE distribution (Cobb et al., 2006) on the interval [2, 6]. u_0 is replaced with an MTE utility function, u_1, by approximating each function $g_1, \ldots, g_4$ with an MTE potential, then combining the result. For example, $g_1(q) = 1/q$ is approximated with an MTE potential $\hat{g}_1(q) = 0.0059 + 0.0452 \cdot \exp(-0.024q)$ by finding constants that minimize the mean squared error (MSE) between g_1 and $\hat{g}_1$ as follows:

$$\underset{a_1, a_2, a_3}{\text{ArgMin}} \int_{q_{min}}^{q_{max}} \left(\frac{1}{q} - a_1 - a_2 \exp(a_3 q) \right)^2 .$$

The values $q_{min} = 63$ and $q_{max} = 110$ are logical lower and upper bounds for Q because values minimize total cost at expected demand at $H = 2$ and $H = 6$. The other functions are similarly approximated.

The expression in (9) for $\hat{S}_R(r)$ can be easily approximated by an MTE function. We find $g_5(r) = -0.0011 + 20.8 \exp(-0.0961r)$ as the function that minimizes the MSE between $\hat{S}_R(r)$ and $g_5(r)$. When the results of the constants in the problem and the approximations to $g_1, \ldots, g_4$ and $\hat{S}_R(r)$ are combined, the result is an MTE function (Moral et al., 2001).

To solve the ID, the variables are deleted in the sequence X, R, Q, H. This example uses the algorithm (Cobb, 2010) that facilitates the determination of a piecewise-linear decision rule for a continuous decision variable with multiple continuous parents. In this case, removing R requires that we determine a decision rule for $R = f(q,h)$. The actual values of Q and H will be known when choosing R.

This decision rule is

$$\hat{\theta}_1(q,h) = \begin{cases} -4.658h - 0.135q + 56.995 & 2.0 \le h < 2.5 \\ -4.658h - 0.129q + 56.462 & 2.5 \le h < 3.0 \\ \vdots & \\ -3.593h - 0.097q + 47.919 & 5.5 \le h \le 6. \end{cases}$$

As Q increases given an established H, the optimal r should be set to a lower value. We divide Ω_H into 8 regions to allow the slope of the piecewise linear decision rule to change for a given value of Q as H increases. We could use more regions to improve accuracy, or fewer regions to speed up the solution process. To complete the elimination of R, we calculate $u_2(q,h) = u_1\left(q, \hat{\theta}_1(q,h), h\right)$. The result is an MTE potential because we are substituting a linear function for R into the MTE potential u_1 (Cobb, 2010).

Applying the decision variable marginalization procedure when a variable has only one continuous parent, as in the case for Q, results in a piecewise linear decision rule as follows: $\hat{\theta}_2(h) = 131.631 - 10.6818h$ for $2 \le h < 3.4$, $\hat{\theta}_2(h) = 126.043 - 9.03846h$ for $3.4 \le h < 4.05$, and $\hat{\theta}_2(h) = 119.18 - 7.34375h$ for $4.05 \le h \le 6$.

Suppose we find before placing an order that $H = 4$. Since $\hat{\theta}_2(4) = 89.9$, we should order $Q = 89.9$ units. Since $\hat{\theta}_1(89.9, 4) = 26.3$, we reorder with $R = 26.3$ units remaining. The expected total cost is 387.0.

5 Discussion

This paper has established a process for estimating the distribution of a sum of i.i.d. log-normal RVs (a single log-normal RV is a special case). The procedure uses the Fenton-Wilkinson approximation (Fenton, 1960) to estimate the parameters for a single log-normal PDF that approximates the sum of log-normal RVs. Once these parameters are determined, an MTE distribution is determined to approximate the PDF with $\mu = 0$ and the FW σ^2 parameter. The method of convolutions is used to find an MTE distribution that approximates that of the sum of the log-normal RVs.

To reduce the tabular inputs required to implement the MTE approximation, the parameters (constants) required to approximate log-normal PDFs with $\mu = 0$ and $0.05 \leq \sigma \leq 1$ are stated as polynomial functions of σ. This covers a range of likely values for the inventory applications presented in the paper, and the FW approximation has been deemed to be effective for these values of σ.

The MTE approach to approximating the distribution for LTD in an inventory control problem was compared to an existing approach to the literature (Das, 1983). One advantage of the MTE method is that a closed-form is obtained for the LTD distribution and the expected shortage per cycle. This advantage was demonstrated when the MTE approximation was subsequently used in an ID model to extend the analysis of the inventory problem to a situation where holding cost is a RV.

Possibilities for future research are as follows. First, the MTE approximation can be extended to cases with larger σ^2 parameters to model applications in such areas as telecommunications and risk analysis. Second, the FW method extends to the case where the component log-normal PDFs are not i.i.d., so the MTE approximation for a sum of log-normals could be utilized in applications with sums of log-normal PDFs that are not identical. Third, other approximations, such as the mixtures of polynomials approach (Li and Shenoy, 2012), could be compared with the MTE approach. Third, extensions of the ID approach to more complicated inventory problems can be pursued.

Acknowledgments

This research has been partially funded by the Spanish Ministry of Economy and Competivity, through project TIN2010-20900-C04-02, by Junta de Andalucía through project TIC-7821 and by ERDF funds.

References

N.C. Beaulieu, A.A. Abu-Dayya, and P.J. McLane. 1995. Estimating the distribution of a sum of independent log-normal random variables. *IEEE Transactions on Communications*, 73(12):2869–2873.

B. R. Cobb, P. P. Shenoy, and R. Rumí. 2006. Approximating probability density functions in hybrid Bayesian networks with mixtures of truncated exponentials. *Statistics and Computing*, 16(3):293–308.

B. R. Cobb. 2010. Continuous decision variables with multiple continuous parents. In P. Myllmäki, T. Roos, and T. Jaakola, editors, *Proceedings of the Fifth European Workshop on Probabilistic Graphical Models*, pages 97–104, Helsinki, Finland.

C. Das. 1983. Inventory control for log-normal demand. *Computers and Operations Research*, 10(3):619–629.

L. F. Fenton. 1960. The sum of log-normal probability distributions in scatter transmission systems. *IRE Transactions on Communication Systems*, 8:57–67.

G. Hadley and T. Whitin. 1963. *Analysis of Inventory Systems*. Prentice-Hall, Englewood Cliffs, N.J.

R. Heuts, K. van Lieshout, and K. Baken. 1986. An inventory model: What is the influence of the shape of the lead time demand distribution? *Mathematical Methods in Operations Research*, 30(2):B1–B14.

H. Langseth, T.D. Nielsen, R. Rumí, and A. Salmerón. 2012. Mixtures of truncated basis functions. *International Journal of Approximate Reasoning*, 53(2):212–227.

Y. Li and P. P. Shenoy. 2012. A framework for solving hybrid influence diagrams containing deterministic conditional distributions. *Decision Analysis*, 9(1):55–75, March.

S. Moral, R. Rumí, and A. Salmerón. 2001. Mixtures of truncated exponentials in hybrid Bayesian networks. In S. Benferhat and P. Besnard, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, Lecture Notes in Artificial Intelligence 2143, pages 156–167. Springer, Berlin.

V. Romero, R. Rumí, and A. A. Salmerón. 2006. Learning hybrid bayesian networks using mixtures of truncated exponentials. *International Journal of Approximate Reasoning*, 42(1-2):54–68.

S.C. Schwartz and Y.S. Yeh. 1982. On the distribution function and moments of power sums with log-normal components. *Bell System Technical Journal*, 61(7):1441–1462.