RC23038 (W0207-114) July 25, 2002
Mathematics

IBM Research Report

A Note on the Inventory Management of High Risk Items

Guillermo Gallego
Columbia University
Department of Industrial Engineering and Operations Research
New York, NY

Kaan Katircioglu, Bala Ramachandran
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

A Note on The Inventory Management of High Risk Items

Guillermo Gallego
Department of Industrial Engineering and Operations Research
Columbia University in the City of New York
Kaan Katircioglu
and
Bala Ramachandran
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598

22 July 2003

Abstract

This note is concerned with the pitfalls of modelling demand as a normal random variable when the coefficient of variation is high. We show that this can lead to considerable under or overestimation of optimal base-stock levels. We propose the lognormal as an alternative model for lead time demands and show that optimal base stock levels asymptotically decrease with the standard deviation. We also present a distribution-free upper bound for optimal base-stock levels, and show that this upper bound first grows linearly with the standard deviation and then remains constant.

1 Problem Statement

We are concerned with the management of inventories in situations where there is a significant uncertainty about demand. Demand uncertainty is often measured by the coefficient of variation cv (defined as the ratio of the standard deviation to the mean). In semiconductor manufacturing, for example, it is common to find parts with $cv > 2$, and it is not unusual to see parts with $cv > 3$.[1] Inventory managers often use the formula $\mu + z\sigma = \mu(1 + zcv)$ with $z \simeq 2$. This formula is based on the normal distribution and results in very high base stock levels, e.g., five times the mean demand, for large cvs. The use of the normal approximation may be even more widespread than intended, as several commercial inventory management packages make the implicit assumption that demand is normal even when the coefficient of variation is large and decisions involve millions of dollars. Some managers defend the normal assumption by stating that they only care about the right tail of the distribution. They feel that the normal distribution is robust against other non-negative distributions with the same mean and variance. Such *faith* in the normal distribution may be misplaced when dealing with large cvs.

[1] In practice, there is often no elaborate use of demand distribution, just guess estimates of the mean and of the forecast error are all one has to calculate safety stocks.

Figure 1 shows actual weekly demand data for a semiconductor product with unit cost \$5.00, selling price \$10.00 and salvage value \$3.00. The empirical distribution coefficient of variation is equal to 2.22, resulting in a sample mean of 207, and a sample standard deviation equal to 459. Although close to three quarters of the demand observations were for fewer than 100 units, there is a chance of receiving a demand for over 1000 units. The Newsvendor solution based on the empirical *cdf* is $Q^* = 100$ resulting in an expected profit of \$63. If we assume demand is normally distributed with the moments calculated based on sample demand data, then the profit maximizing solution will be 467 units resulting in an expected *loss* of \$291 (based on the empirical distribution). To satisfy demand with probability 95%, management would have to order 1,400 units and incur a *loss* of \$1,583.

This example illustrates the danger of using the normal distribution when the coefficient of variation is high. This is consistent with the findings of other authors. See Bookbinder and Lordahl [2] and Lordahl and Bookbinder [3] for examples where alternative methods for calculating inventory policies perform better than solutions based on the normality assumption. Moreover, empirical studies, see Agrawal and Simth [1], suggest that skewed distributions, such as the Negative Binomial, may provide a better fit than the normal for situations where unmet demand is lost.

The Gamma distribution is the continuous counterpart of the Negative Binomial distribution; the scale and location parameters of the Gamma admit all possible non-negative means and positive standard deviations. Unfortunately, while it is possible to write partial expectation of the Gamma in terms of the incomplete Gamma function, the need to integrate the incomplete Gamma function makes it numerically expensive to perform computations, particularly for large coefficient of variations.

In this note we propose the lognormal as an alternative to the normal to manage high risk items. The reasons for selecting the lognormal are based on observed skewed plots, the ease of computing base stock levels and the corresponding expected costs. As we shall see, the lognormal leads to more conservative base stock levels for large *cv*s. We envision that our results will be useful, not only to suggest sensible policies for high risk items when the distributional assumptions are justified, but also for answering what if questions related to the financial impact of changes in distributional parameters. We believe that in the absence of detailed distributional information the lognormal provides a good option as a default demand model for high risk items. We finalize this note by providing an upper bound on the base stock levels for all non-negative distributions with a given mean and variance.

2 Inventory Models Minimizing Costs

Most of the issues concerning large *cvs*, e.g., large inventory at risk, are more serious for products with short life cycles. To simplify the presentation of the relevant issues, we model the case where there is a single opportunity to order and demand occurs over a single period. We assume that demand D is a non-negative random variable with mean μ and variance σ^2. To simplify the exposition we will also assume that we start with zero inventories. Extensions to positive initial inventories and multi-period models with positive lead times are straightforward. See Veinott [7] and references therein.

Let c be the unit cost, $p > c$ the unit selling price, and $s < c$ the unit salvage price. In addition, let g be a goodwill cost per unit of unsatisfied demand. If we decide to be in business, the expected profit of ordering $y > 0$ units is given by $\pi(y) = (p - c)\mu - G(y)$, where $G(y) = HE(y - D)^+ + BE(D - y)^+$, $H = c - s$ and $B = p-c+g$. Thus, maximizing the expected profit is equivalent to minimizing $G(y)$, over $y > 0$. Minimizing G is a Newsvendor problem with overage cost H and underage cost B.[2] Under well known conditions, G exists, is convex, and has a finite minimizer y^*, e.g., the smallest y such that $P(D \leq y) \geq \beta \equiv B/(B + H)$, so y^* satisfies all the demand with probability at least $100\beta\%$. Because the Newsvendor problem is often cast as a cost minimization problem, the analysis usually stops after y^* is found. However, $\pi(y^*)$ may be negative when the Newsvendor costs $G(y^*)$ is larger than the profit contribution $(p - c)\mu$. This can only happen if $g > 0$. Since it is reasonable to assume that goodwill costs do not accrue if we are not in business, the expected profit and the optimal policy should be modified so that no order is placed if ordering leads to an expected loss. In our experience, the sign of $\pi(y^*)$ is seldom checked because the possibility is not included in some commercial inventory packages and more often because a service model is used instead of a cost model.

2.1 The Normal Approximation

Under the normal approximation, $y^* = \mu+\sigma z_\beta$ where $z_\beta = \Phi^{-1}(\beta)$. Then $\pi(y*) = (p-c)\mu-(H+B)\sigma\ \phi(z_\beta)$.

2.2 The Maximal Approximation

The maximal approximation minimizes cost against the worst *non-negative* distribution with a given mean and a given variance. See Scarf [5] and Gallego and Moon [4]. The base stock level under the maximal approximation is given by $y^* = \mu + \frac{1}{2}\sigma\left(\sqrt{\frac{B}{H}} - \sqrt{\frac{H}{B}}\right)$ if $cv < \sqrt{B/H}$ resulting in the expected profit lower bound $\pi(y^*) = (p - c)\mu - \sigma\sqrt{HB}$. If $cv > \sqrt{B/H}$, it is optimal *not* to order under the maximal

[2] H and B are also known as the holding and backorder cost rates when dealing with multi-period problems.

approximation.

2.3 The Lognormal Approximation

A random variable D is said to have the lognormal distribution, with parameters ν and τ, if $\ln(D)$ has the normal distribution with mean ν and standard deviation $\tau \geq 0$. The lognormal distribution is often used to model non-negative random variables such as lifetimes and total returns. It is well known that $E(X^n) = \exp(n\nu + n^2\tau^2/2)$. Thus, $\mu = \exp(\nu + \tau^2/2)$ and $\sigma^2 = \mu^2(\exp(\tau^2) - 1)$, so $\nu = \ln\mu - \ln\sqrt{1 + cv^2}$ and $\tau = \sqrt{\ln(1 + cv^2)}$.

The following Proposition gives the base-stock level and the expected profit for the lognormal distribution.

Proposition: 1 $y^* = \exp(v + \tau z_\beta)$ *and* $\pi(y^*) = (p - c)\mu - (H + B)\mu\Phi(\tau - z_\beta) + H\mu.$

Proof of Proposition 1

If D is lognormal then $Pr(D \leq y^*) = Pr(\ln(D) \leq \ln(y^*)) = Pr(\nu + \tau Z \leq \nu + \tau z_\beta) = \Phi(z_\beta) = \beta$. Now, using the fact that $E(D - y^*)^+ = \mu\Phi(\tau - z_\beta) - y^*\Phi(-z_\beta)$ and $\Phi(-z_\beta) = H/(H + B)$ we see that

$$\begin{aligned} G(y^*) &= H(y^* - \mu) + (H + B)E(D - y^*)^+ \\ &= H(y^* - \mu) + (H + B)\mu\Phi(\tau - z_\beta) - (H + B)y^*\Phi(-z_\beta) \\ &= (H + B)\mu\Phi(\tau - z_\beta) - H\mu. \end{aligned}$$

2.4 Profit Comparison

We now compare the performance of base-stock levels set by the normal, maximal and the lognormal formulas in terms of the expected profit. Managers who use the normal model to compute base stock levels intend to use the resulting base stock levels against a non-negative distribution with the same mean and variance. For this reason, we test these base-stock levels against the worst case, the lognormal and the gamma distributions. The shape and scale parameters of the gamma distributions are selected to match the first two moments of the distribution. For comparison purposes, we also test the performance against the normal distribution. Test against the normal are only relevant for small coefficients of variations, e.g., $cv \leq 1/3$. For large cv's the results are distorted by the mass assigned to negative values. For this study we use the following parameters: $c = 100$, $p = 200$, $s = 25$, $g = 10$, $\mu = 100$ and $\sigma \in \{30, 100, 200, 300\}$. Table 1 contains the results. As we can see, using the normal can lead to significant losses. The lognormal on the other hand, performs better against the gamma and the worst case distribution.

True Distribution	y^*	NORMAL	Worst Case	LOGNORMAL	GAMMA
$c = 100, p = 200, s = 25, g = 10, \mu = 100, \sigma = 30$					
NORMAL IP	107	\$7,848.41	\$7,272.27	\$7,827.67	\$7,813.47
MAXIMAL IP	106	7,846.05	7,275.11	7,839.65	7,803.60
LOGNORMAL IP	103	7,824.76	7,261.54	7,850.32	7,824.11
$c = 100, p = 200, s = 25, g = 10, \mu = 100, \sigma = 100$					
NORMAL IP	124	NOT REL.	\$907.58	\$3,525.03	\$2,901.95
MAXIMAL IP	119	NOT REL.	917.05	3,652.85	2,999.13
LOGNORMAL IP	86	NOT REL.	510.35	4,115.85	3,307.80
$c = 100, p = 200, s = 25, g = 10, \mu = 100, \sigma = 200$					
NORMAL IP	148	NOT REL.	-\$6,619.29	-\$473.82	-\$2,656,02
MAXIMAL IP	0	NOT REL.	0	0	0
LOGNORMAL IP	61	NOT REL.	-3,302.44	1,806.07	-\$114.53
$c = 100, p = 200, s = 25, g = 10, \mu = 100, \sigma = 300$					
NORMAL IP	172	NOT REL.	-\$10,707.50	-\$3,318.67	-\$3,477.10
MAXIMAL IP	0	NOT REL.	0	0	0
LOGNORMAL IP	45	NOT REL.	-3,569.23	861.41	\$127.34

Table 1: Expected Profit Comparison (IP- Inventory Policy)

2.5 Semiconductor Example Revisited

If we use lognormal distribution with the sample moments, the profit maximizing solution will be 181 units giving us an expected profit of \$29. The maximal distribution will abstain from ordering. Our purpose here is not to test the fitness of lognormal but to illustrate how the lognormal ordering policy may result in a more sensible solution in a real world example.

2.6 Skewness of Lognormal

We end this section with a result that explains the asymptotic behavior of the lognormal as σ increases. In particular, in sharp contrast to the normal, it follows that base stock levels tend to zero when σ is large.

Theorem 1 *If D is lognormal with mean μ and variance σ^2 then for all $x > 0$,*

$$Pr(D \leq x) \to 1 \quad \text{as } \sigma \to \infty. \tag{1}$$

Proof: Since $\ln(\sqrt{1+cv^2})/\sqrt{\ln(1+cv^2)} \to \infty$ as $\sigma^2 \to \infty$, it follows that

$$Pr(D \leq x) \quad = \quad \Phi\left(\frac{\ln(\frac{x}{\mu}) + \ln(\sqrt{1+cv^2})}{\sqrt{\ln(1+cv^2)}}\right) \to 1$$

for all $x > 0$.

This result tells us that for fixed μ, the probability that a lifetime ends by time x, for *any* $x > 0$, becomes certain as $\sigma \to \infty$. In sharp contrast, for the normal distribution, the probability that a lifetime ends by time x *decreases* with σ for *all* $x > \mu$. The convergence of equation (1) is not monotone for values of $x > \mu$. Indeed, for $x > \mu$, the probability $P(D \leq x)$ decreases and then increases with σ.

3 Inventory Models Minimizing Costs with Service Constraints

Managers often use service levels instead of a cost model to determine base-stock policies by selecting a desired service level β. The normal prescribes a base stock level $y_n(\beta) = \mu + \sigma z_\beta$ to satisfy demand with probability $100\beta\%$. In contrast, the lognormal model prescribes a base stock level $y_{ln}(\beta) = \exp(v + \tau z_\beta)$. Notice that $y_{ln}(\beta) \to 0$ as $\sigma \to \infty$. It is easy to see that $y_{ln}(\beta)$ is maximized when $cv^2 = \exp(z_\beta^2) - 1$, so $y_{ln}(\beta) \leq \mu \exp\left(\frac{1}{2} z_\beta^2\right)$.

It is instructive to compare how the base stock levels for the normal and the lognormal compare. At low cvs, the lognormal may prescribe a higher base stock level than the normal distribution, but the normal prescribes a larger base-stock level beyond the crossover point, defined as the largest cv, say cv_β, such that $y_{ln}(\beta) = y_n(\beta)$. Table 2 gives the ratio $y(\beta)/E[D]$ as a function of the coefficient of variation cv, for values of $\beta \in \{.75, .9, .98\}$. It also provides the cv that maximizes the lognormal base stock level, and the crossover point with the normal. Notice, for example, that for $\beta = 0.75$, the normal prescribes a base stock level that is more than three times as large than the lognormal distribution when $cv = 3$.

$z_{.75} = 0.674$, $y_{ln}(0.75)$ peaks at $cv = 0.76$; $cv_{0.75} = 0.04$								
cv	1	2	3	4	5	10	25	50
$\frac{y_{ln}(.75)}{E[D]}$	1.24	1.05	0.88	0.75	0.66	0.42	0.22	0.13
$\frac{y_n(.75)}{E[D]}$	1.67	2.35	3.02	3.70	4.37	7.74	17.86	34.72
$z_{.90} = 1.28$, $y_{ln}(.90)$ peaks at $cv = 2.04$; $cv_{.90} = 0.49$								
cv	1	2	3	4	5	10	25	50
$\frac{y_{ln}(.90)}{E[D]}$	2.06	2.27	2.21	2.10	1.98	1.56	1.03	0.72
$\frac{y_n(.90)}{E[D]}$	2.28	3.56	4.84	6.13	7.41	13.82	33.04	65.08
$z_{.98} = 2.05$, $y_{ln}(.98)$ peaks at $cv = 8.12$; $cv_\beta = 2.98$								
cv	1	2	3	4	5	10	25	50
$\frac{y_{ln}(.98)}{E[D]}$	3.91	6.05	7.14	7.69	7.99	8.20	7.33	6.25
$\frac{y_n(.98)}{E[D]}$	3.05	5.11	7.16	9.21	11.27	21.54	52.34	103.69

Table 2: Relative Base-Stock Levels

3.1 Performance of Base Stock Levels Against other Distributions

We now compare the performance of base-stock levels set by the normal and the lognormal formulas. Table 3 shows the results for the case with lead time demand 100, $cv = 2$ for $\beta \in \{0.75, 0.90, 0.98\}$. It is clear that the base-stock policy generated by the normal distribution grossly over-estimates optimal base-stock levels for $\beta = 0.75$, slightly over-estimates optimal base-stock levels for $\beta = 0.90$, and underestimates optimal base-stock levels for $\beta = 0.98$. On the other hand, the base-stock policy generated by the lognormal distribution is on target relative to the gamma for $\beta = 0.75$ and $\beta = 0.98$, but underestimates the base stock for gamma demand for $\beta = 0.90$. Table 4 provides the results for $cv = 3$. It is clear from the non-negative distribution results from Tables 3 & 4 that the lognormal does not suffer from the overestimation of base stock policy for low service levels, while also performing better than the normal for achieving high service levels.

$\beta = 0.75$; $y_n(.75) = 235$, $y_{ln}(.75) = 105$.

Demand	LogNormal	Gamma
LognormalPolicy	75%	75%
NormalPolicy	90%	87%

$\beta = 0.90$; $y_n(0.90) = 356$, $y_{ln}(.90) = 227$.

Demand	LogNormal	Gamma
LognormalPolicy	90%	86%
NormalPolicy	95%	92%

$\beta = 0.98$; $y_n(\beta) = 511$, $y_{ln}(.98) = 605$.

Demand	LogNormal	Gamma
LognormalPolicy	98%	97%
NormalPolicy	96%	95%

Table 3: Performance of Policies for $cv = 2$

3.2 Non-Negative Distributions with Finite First and Second Moments

In this section we provide a tight upper bound on base stock levels for non-negative random variables with mean μ and variance σ^2. Let D be a non-negative random variables with mean μ and variance σ^2 and let $y(\beta)$ be the smallest base stock level y such that $P(D \geq y) \leq 1 - \beta$. Let

$$y^+(\beta) = \begin{cases} \mu + \sqrt{\frac{\beta}{1-\beta}}\sigma & \text{if } \sigma < \mu\sqrt{\frac{\beta}{1-\beta}} \\ \frac{\mu}{1-\beta} & \text{otherwise.} \end{cases}$$

Theorem 2

$$y(\beta) \leq y^+(\beta) \tag{2}$$

Proof of Theorem 2

$\beta = 0.75$; $y_n(.75) = 302$, $y_{ln}(.75) = 88$.

$Demand$	LogNormal	Gamma
$Lognormal Policy$	75%	81%
$Normal Policy$	93%	91%

$\beta = 0.90$; $y_n(0.90) = 484$, $y_{ln}(.90) = 221$.

$Demand$	LogNormal	Gamma
$Lognormal Policy$	90%	88%
$Normal Policy$	96%	94%

$\beta = 0.98$; $y_n(\beta) = 715$, $y_{ln}(.98) = 708$.

$Demand$	LogNormal	Gamma
$Lognormal Policy$	98%	96%
$Normal Policy$	98%	96%

Table 4: Performance of Policies for $cv = 3$

In this section we investigate the behavior of $P(D < t)$ for fixed $t > \mu$ for a non-negative random variable with mean μ and variance σ^2. Part of the answer is given by Cantelli's inequality:

$$Pr(D < t) \geq \frac{(t-\mu)^2}{\sigma^2 + (t-\mu)^2}. \tag{3}$$

It can be easily shown that Cantelli's inequality, see Stoyan [6], is tight for non-negative distributions if and only if $t \geq E[D^2]/\mu$. When $\mu \leq t < E[D^2]/\mu$ we can improve on Cantelli's inequality by using the restriction that D is non-negative. Markov's inequality tells us that

$$Pr(D < t) \geq \frac{t-\mu}{t} \tag{4}$$

for all nonnegative random variables D with mean μ. Markov's inequality is achieved by the two point distribution with mass at 0 and t and probabilities $(t-\mu)/t$ and μ/t. This distribution has variance $(t-\mu)\mu < \sigma^2$ for all $\mu \leq t < E[D^2]/\mu$. Thus, although Markov's inequality provides an improvement over Cantelli's inequality for $\mu \leq t < E[D^2]/\mu$, it is not clear whether a better bound can be found by restricting to distributions with variance equal to σ^2. Lemma 1 shows that this is not possible. In other words, Markov's inequality is tight over this region.

From Cantelli's inequality, we have $y_c(\beta) = \mu + \sigma\sqrt{\beta/(1-\beta)}$. For Markov's inequality, we have $y_m(\beta) = \mu/(1-\beta)$. Cantelli's inequality is tight when $\sigma \leq \mu\sqrt{\beta/(1-\beta)}$. Markov's inequality is tight otherwise. Combining the two bounds results in $y^+(\beta)$.

Lemma 1 *Markov's inequality (4) is tight for all non-negative random variables D with mean μ and variance σ^2 for $\mu \leq t < E[D^2]/\mu$.*

Proof: Let $t' > t$ and let $\alpha(t') = (t-\mu)t'/(t'-\mu)t \in [0, 1)$. We will construct a three point distribution with mass at $0, t$, and t' and respective weights $(t-\mu)/t$, $\mu\bar{\alpha}(t')/t$ and $\mu\alpha(t')/t'$, where $\bar{\alpha}(t') \equiv 1 - \alpha(t')$. It is easy to see that this distribution has mean μ and $P(D < t) = (t - \mu)/t$. Moreover, the variance of this distribution,

$$\sigma^2(t') \equiv \mu(t - \mu)/t + \alpha(t')\mu(t' - \mu) + \bar{\alpha}(t')(t - \mu)^2\mu/t$$

increases without bound as t' increases. Therefore, there exists a $t' \geq t$ such that $\sigma^2(t') = \sigma^2$. We have shown that there exists a nonnegative distribution with mean μ and variance σ^2 that achieves Markov's lower bound. Thus, Markov's lower bound is tight for distributions with $\sigma^2 \geq (t - \mu)\mu$.

$\square$

Notice that for fixed μ and β, $y^+(\beta)$ initially increases linearly with σ, but is a constant for large σ. This result can be used, for example, to curb the normal distribution from generating inappropriately large base stock levels, by setting $y(\beta) = \min(y^+(\beta), y_n(\beta))$.

Let $F(x) = P(D \leq x)$. For continuous distributions $F(y(\beta)) = \beta$, or equivalently $F(x) = y^{-1}(x)$. One may question if there is a continuous distribution that achieves $y^+(\beta)$ for all $\beta \in [0, 1)$. This is, indeed the case, and the distribution is given by inverting $y^+(\beta)$. This results in $F(x) = 0$ for $0 \leq x < \mu$, $F(x) = 1 - \mu/x$ for $\mu \leq x < \frac{\sigma^2+\mu^2}{\mu}$ and $F(x) = \frac{(x-\mu)^2}{\sigma^2+(x-\mu)^2}$ for $x > \frac{\sigma^2+\mu^2}{\mu}$. Notice, by integrating the tail of the distribution, that this distribution does not have the prescribed moments.

4 Managerial Implications

Practical wisdom suggests that larger safety stocks are needed to protect against demand uncertainty when the standard deviation of demand increases. A key message of this paper is that this is not always true. Negative safety stocks, i.e., base stock levels below the mean, can result in high service levels. In these cases, it may be enormously expensive and unnecessary to insist on a base stock level of the form $\mu + z\sigma$ as suggested by the normal distribution. In addition, there is a threshold above which the base stock level becomes insensitive to changes in the standard deviation. This follows from our bound on base stock levels that first grows with σ and then remains constant. Curbing base stock levels using this upper bound may avoid excessively large orders suggested by the use of inappropriate distributions.

References

[1] Agrawal, N. and S. Smith. 1996. Estimating Negative Binomial Demand for Retail Inventory Management with Unobservable Lost Sales. *Naval Research Logistics*, 43, 839-861.

[2] Bookbinder, J.H. and Lordahl, A.E. 1994. Order-Statistics Calculation, Costs and Service in an (s,Q) Inventory System, *Naval Research Logistics*, 41, 81-97.

[3] Lordahl, A.E. amd Bookbinder, J.H. 1988. Estimation of Inventory Re-order Levels using the Bootstrap Statistical Procedure, *IIE Transactions*, 21, 4.

[4] Gallego, G., and Moon, I. 1993. The Distribution Free Newsboy Problem: Review and Extensions. *Journal of the Operational Research Society.* 44, 825-834.

[5] Scarf, H. 1958. A Min-Max Solution of an Inventory Problem, *Ch. 12 in Studies in The Mathematical Theory of Inventory and Production, Stanford Univ. Press.*

[6] Stoyan, D. 1983. *Comparison Methods For Queues and Other Stochastic Models*, John Wiley, NY.

[7] Veinott, A. 1965. Optimal Policy for a Multi-Product, Dynamic, Nonstationary Inventory Problem. *Management Science*, 12, 3, 206-222.

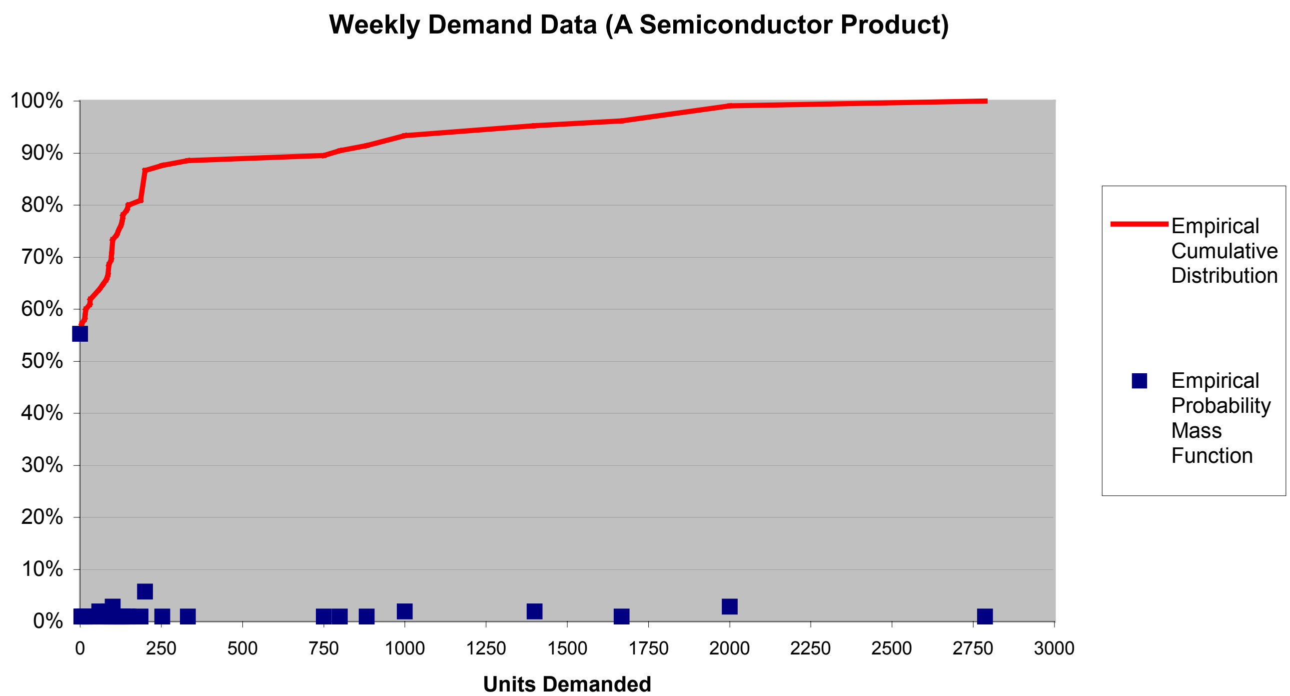

Figure 1: Weekly Demand Data for A Semiconductor Product (Mean=207, Stdev=459.)