Inventory control with stochastic lead times: A tutorial

Stephen M. Disney

UEBS-ASU Interdisciplinary Workshop on Supply Chain Resilience,
June 11th-13th, Exeter, UK.

and

Recent Advances in Supply Chain Forecasting:
A Workshop in Memory of Professor John E. Boylan,
June 13th-14th, 2024, Lancaster University Business School, UK.

Global supply chains have plenty of opportunities for random delays and hold-ups ♠

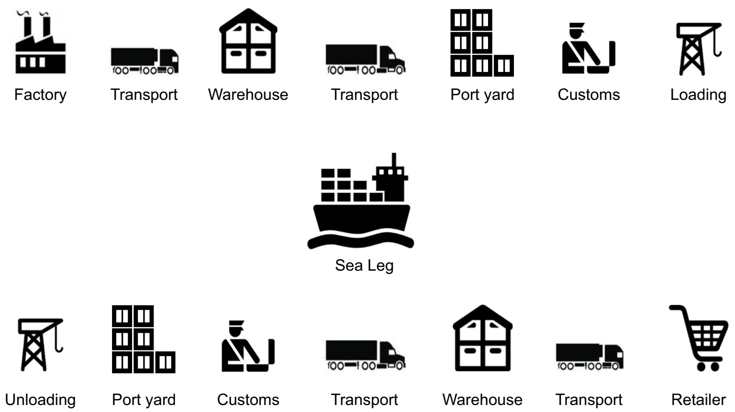

Longshoreman/stevedore strikes, regulations & slow custom checks, lack of truck drivers, piracy, canal blockages, bridge crashes; the list is endless.

Overview

- Distribution of stochastic lead times.
- Order cross-overs.
- Random sums of random numbers.
- EOQ safety stocks with stochastic lead times.
- Expected dynamic behavior of the POUT policy with stochastic lead times drawn from a negative binomial distribution.
- POUT policy with stochastic lead times: When an arbitrary lead time distribution is based on empirical evidence (a posteriori).
- Open orders and sub-processes.
- The inventory and order distributions.
- OUT policy with stochastic lead times: Forecasting the lead time.
- The case of auto-correlated demand and correlated lead times.

Distribution of lead time for port-to-port shipping from China to the USA during March-September 2011

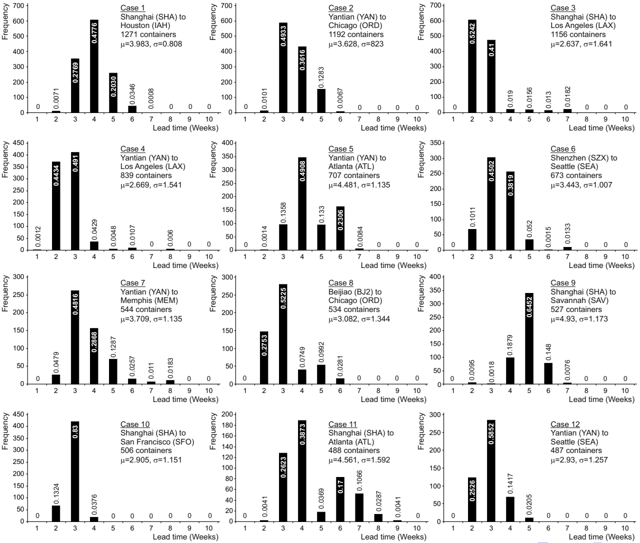

Factory to factory lead times between Boulder and Shenzhen are stochastic and have order crossovers

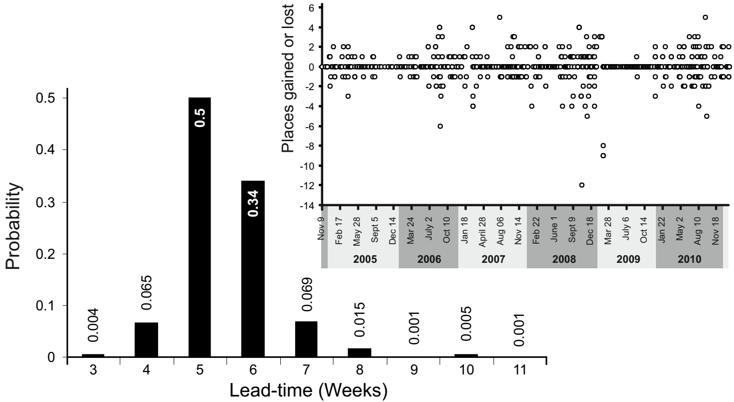

Orders that **crossover** arrive in a different sequence than they were sent.

Sum of random variables

- Assume demand D, is an i.i.d. normally distributed random variable with a probability density function (pdf) of

$$f_D = e^{-\frac{(x-\mu_D)^2}{2\sigma_D^2}} \Big/ \sqrt{2\pi\sigma_D^2} \tag{1}$$

- Using the convolution integral, the pdf of the demand over the lead-time L is given by the L-fold convolution,

$$\phi_{\sum_{i=1}^{L} D_i}[x] = (f_{D_1} \star f_{D_2} \star \cdots f_{D_L}) = e^{-\frac{(x-L\mu_D)^2}{2L\sigma^2}} \Big/ \sqrt{2\pi\sigma_D^2 L}, \tag{2}$$

which is the pdf of a normal distribution with a mean of $L\mu_D$ and a variance of $L\sigma_D^2$.

- Put simply, the pdf of a sum i.i.d. random variables has a mean equal to the sum of the means and a variance equal to the sum of the variances.

Safety stocks with constant lead times: A numerical example

- Assume a normally distributed demand has a mean of $\mu_D = 100$ and variance of $\sigma_D^2 = 100$.
- 95% availability is desired, hence a safety factor of $\Phi^{-1}[0.95] = 1.645$ is used.[1]
- With a constant lead time of $L = 4$ the following re-order point RP is set

$$RP = \left\lceil L\mu_D + \Phi^{-1}[0.95]\sqrt{L\sigma_D^2} \right\rceil = \left\lceil 4(100) + 1.645\sqrt{4(100)} \right\rceil = 433 \quad (3)$$

[1] The Shiny app at https://bullwhip.shinyapps.io/StatisticalDistributions/ provides more information on the normal distribution.

Safety stocks with random lead times, Eppen and Martin (1988) ♠

- Assume lead time $L = 2$ with probability $p_2 = 0.5$ and $L = 4$ with probability $p_4 = 0.5$.
- Mean $\mu_L = \sum_{L=1}^{L_+} p_L L = 3$ and variance $\sigma_L^2 = \sum_{L=1}^{L_+} p_L (L - \mu_L)^2 = 1$.
- Normally distributed demand has a mean of $\mu_D = 100$ and variance of $\sigma_D^2 = 100$.
- Distribution of the demand over the lead-time is the sum of a random number of random variables whose pdf is given by

$$\tilde{\phi}[x] = \sum_{L=1}^{L_+} p_L \phi_{\sum_{i=1}^{L} D_i}[x]. \quad (4)$$

  Here, p_L is the probability of lead-time L occurring and L_+ is the maximum lead-time, Feller (1958).
- The lead-time demand has a mean of $\mu_D \mu_L$ and a variance of $\sigma_D^2 \mu_L + \sigma_L^2 \mu_D^2$, Feller (1958).
- This mean and variance **incorrectly** implies the re-order point should be set to

$$RP = \left\lceil \mu_D \mu_L + \Phi^{-1}[0.95]\sqrt{\sigma_D^2 \mu_L + \sigma_L^2 \mu_D^2} \right\rceil = \left\lceil 3(100) + 1.645\sqrt{10300} \right\rceil = 467. \quad (5)$$

Feller's formula gives correct mean & variance of lead-time demand but is not sufficient to describe the complete pdf

The true pdf of the lead-time demand is given by $\tilde{\Phi}_{\sum_{i=1}^{L} D_i}[x] = \frac{1}{2}\sum_{L=1}^{L_+} p_L(1 + erf[\frac{x-L\mu}{\sqrt{2}\sigma}])$

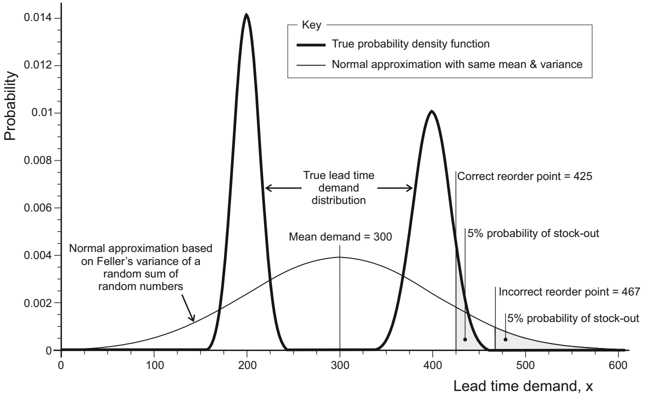

Which provides the re-order point, RP

$$\tilde{\Phi}_{\sum_{i=1}^{L} D_i}[x] = \frac{1}{2} + \frac{1}{4}\text{erf}[\frac{RP}{20} - 10] + \frac{1}{4}\text{erf}[\frac{RP-400}{20\sqrt{2}}] = 0.95 \implies RP = 425$$

The random sum of random numbers approach should be avoided

- When the modes of a multi-modal inventory distribution are sufficiently far apart, the reorder point calculation may be dominated by the greatest mode.
- If the right-hand tails of the pdf's are similar, then the two re-order points can be close. This happens when the distribution of the lead times is bell-shaped and when the variance of the demand is rather large compared to the mean demand.
- Neither-the-less, even if the results are similar from both approaches, we recommend that the multi-modal approach is taken as;
  - it is hard to predict when a good approximation is present,
  - is not difficult to calculate in Excel,
  - is exact,
  - and is correct for the correct reason.

Dynamic response of the proportional order-up-to (POUT) policy

- Consider an impulse demand is present. The impulse demand is $D_t = 1$ when time $t = 0$ and $D_t = 0$ for all other t.
- The impulse demand can be used to represent i.i.d. demand as the demand impulse response is also a system's auto-covariance function.
- The minimum mean squared error (MMSE) forecast of i.i.d. demand is $\hat{D}_{t+j|t} = \mu_D$, $\forall j$.
- Balakrishnan et al. (2004) and Boute and Van Mieghem (2015) show the POUT policy under i.i.d. demand with MMSE forecasts generates replenishment orders via

$$O_t = \beta D_t + (1 - \beta) O_{t-1} \tag{6}$$

- The inventory balance equation is

$$I_t = I_{t-1} + R_t - D_t. \tag{7}$$

  Here, I_t is the inventory at time t, and R_t is the production completions (or suppliers' receipts).
- We can model the lead-time, the link between O_t and R_t, in many different ways.

Production/replenishment lead time distributions ♠

- **First-order exponential smoothing**. Parts of the order arrive over many periods in the future. Maybe representative of factory output where production is completed over several periods, some output arrives immediately, and some output is spread out over time.
- **Third-order exponential smoothing**. Equivalent to three first-order exponential smoothing mechanisms in series. Representative output following an s-shaped learning curve, Naim (1993).
- **Negative binomial distribution**. A generalization of the n^{th}-order exponential smoothing models.
- **Geometric distribution**. A first-order negative binomial distribution (first-order exponential smoothing) is equivalent to a geometric distribution.
- **Pure time delay**. When $n \to \infty$, the negative binomial distribution models a pure time delay. Suitable when only complete orders arrive in future periods. Representative of a transportation delay.

Expected dynamic behaviour

- The impulse response of an n^{th}-order exponential smoothing delay can be interpreted as the **expected dynamic behavior** of a system with a stochastic delay drawn from a negative binomial distribution.
- The connection is established by observing that r_t, the inverse z-transform of the impulse response of the n^{th} order exponential smoothing (with smoothing parameter α), is identical to the probability mass function of the negative binomial distribution of the number of failures k, with p probability of success in each trial until r successes have occurred.

$$f(k, r, p) = \binom{k+r-1}{k} p^r (1-p)^k \qquad \text{(Negative binomial distribution)}$$

$$R_t = \binom{t+n-1}{t} \alpha^n (1-\alpha)^t \qquad (n^{th}\text{-order exponential smoothing})$$

- Mapping $k \to t$, $r \to n$, and $p \to \alpha$, reveals that we can use the pmf of the negative binomial distribution to represent the pmf of lead times.
- When we do so, the expected dynamic response can be obtained by replacing the stochastic lead times with an n^{th}-order exponential smoothing model.

Expected dynamic response: The impulse response of the POUT policy's inventory levels ($\beta = 0.3$)

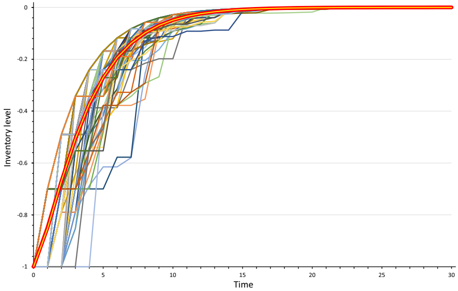

- Figure shows 100 sample paths of the inventory impulse response with a random lead time. The lead time distribution is an **geometric distribution** (a first-order negative binomial distribution, $r = 1$), with parameter $p = 0.5$.
- The red line is the average dynamic response of 10000 sample paths.
- The yellow line is the inventory impulse response with a **first-order exponential smoothing** (with $\alpha = 0.5$) representing the lead time.
- Upper boundary of the sample paths is the impulse response with zero lead time.

The proportional order-up-to (POUT) policy with a known stochastic lead time distribution, Disney et al. (2016)

- Consider a normally distributed, i.i.d. random demand (D_t), with mean μ_D and variance σ_D^2, is present.
- The minimum mean squared error forecast of i.i.d. demand is $\hat{D}_{t+j|t} = \mu_D$, $\forall j$.
- The inventory balance equation is

$$I_t = I_{t-1} + O_{t-1-L} - D_t. \tag{8}$$

Here: I_t is the inventory at time t, L is a discrete i.i.d. random variable that represents the lead time of the order O_{t-1-L}.

- The POUT policy generates replenishment orders via

$$O_t = \mu_D + \beta(T + L\mu_D - (I_t + W_t)). \tag{9}$$

Here: T is a target safety stock. L is the lead time. W_t is the inventory on order, the work-in-progress (WIP), the orders placed but not received.

The impact of the proportional feedback controller, β, under i.i.d. demand and constant lead times ♠

- β regulates how quickly deviations in inventory and WIP levels are corrected.
- $0 \leq \beta < 2$ is required for stable operation.
- When $\beta = 1$, the POUT policy degenerates into the OUT policy.
- When a **constant** lead time L is present the order variance is given by,

$$\mathbb{V}[O_t] = \frac{\sigma_D^2 \beta}{2 - \beta} \tag{10}$$

and the inventory variance is

$$\mathbb{V}[I_t] = \sigma_D^2 \left(L + \frac{1}{\beta(2 - \beta)} \right). \tag{11}$$

The inventory and order variance trade-off in the POUT policy under constant lead times

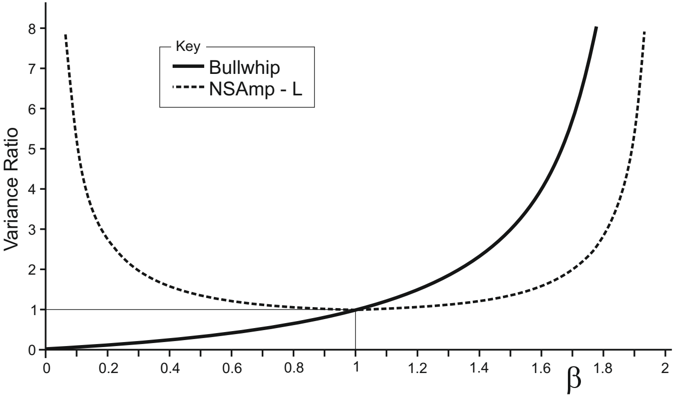

Note: Bullwhip= $\frac{\mathbb{V}[O_t]}{\mathbb{V}[D_t]}$ and NSAmp = $\frac{\mathbb{V}[I_t]}{\mathbb{V}[D_t]}$.

The state of the WIP pipeline ♠

- The key to understanding the impact of the stochastic lead time is to consider the number of **open** orders in the WIP pipeline.
- Open orders are those that have been placed but not yet received. Notice, we are not describing the **quantity** of products on order, but the **number** of open orders.
- Let the probability of an order having a lead time of L be given by p_L.
- L^+ is the maximum lead time; the average lead time is $\mu_L = \sum_{L=0}^{L^+} p_L L$.
- All orders placed L^+ or more periods ago have been received.
- However, orders placed later than $L^+ - 1$ periods ago may be open (not yet received, denoted by '1') or closed (received, denoted by '0').
- Since each of the L^+ positions in the pipeline is either open or closed, this means that there are 2^{L^+} possible *states* of the WIP pipeline.

Probability of the pipeline being in state i

- The probability that the pipeline is in state i is denoted q_i.
- The relationship between p_L and q_i is rather complex.
- To explain the relationship, consider a case where the lead time possibilities are $p_0 = 0$, $p_1 = 1/3$, $p_2 = 1/2$ and $p_3 = 1/6$.
- Note that the probabilities sum to unity and the maximum lead time is $L^+ = 3$.

The combinations of open orders and their associated probabilities, $p_0 = 0$, $p_1 = 1/3$, $p_2 = 1/2$ and $p_3 = 1/6$

State i	$t-1, j=1$	$t-2, j=2$	$t-3, j=3$	$t-4, j=4$	Probability, $q_i = \prod_{j=1}^{4} q_{i,j}$
1	0 $q_{1,1} = 0$	0 $q_{1,2} = 0 + \frac{1}{3}$	0 $q_{1,3} = 0 + \frac{1}{3} + \frac{1}{2}$	0 $q_{1,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_1 = 0$
2	0 $q_{2,1} = 0$	0 $q_{2,2} = 0 + \frac{1}{3}$	1 $q_{2,3} = 1 - (0 + \frac{1}{3} + \frac{1}{2})$	0 $q_{2,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_2 = 0$
3	0 $q_{3,1} = 0$	1 $q_{3,2} = 1 - (0 + \frac{1}{3})$	0 $q_{3,3} = 0 + \frac{1}{3} + \frac{1}{2}$	0 $q_{3,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_3 = 0$
4	0 $q_{4,1} = 0$	1 $q_{4,2} = 1 - (0 + \frac{1}{3})$	1 $q_{4,3} = 1 - (0 + \frac{1}{3} + \frac{1}{2})$	0 $q_{4,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_4 = 0$
5	1 $q_{5,1} = 1 - 0$	0 $q_{5,2} = 0 + \frac{1}{3}$	0 $q_{5,3} = 0 + \frac{1}{3} + \frac{1}{2}$	0 $q_{5,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_5 = \frac{5}{18}$
6	1 $q_{6,1} = 1 - 0$	0 $q_{6,2} = 0 + \frac{1}{3}$	1 $q_{6,3} = 1 - (0 + \frac{1}{3} + \frac{1}{2})$	0 $q_{6,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_6 = \frac{1}{18}$
7	1 $q_{7,1} = 1 - 0$	1 $q_{7,2} = 1 - (0 + \frac{1}{3})$	0 $q_{7,3} = 0 + \frac{1}{3} + \frac{1}{2}$	0 $q_{7,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_7 = \frac{5}{9}$
8	1 $q_{8,1} = 1 - 0$	1 $1 - (q_{8,2} = 0 + \frac{1}{3})$	1 $q_{8,3} = 1 - (0 + \frac{1}{3} + \frac{1}{2})$	0 $q_{8,4} = 0 + \frac{1}{3} + \frac{1}{2} + \frac{1}{6}$	$q_8 = \frac{1}{9}$

0: Closed, received orders. 1: Open, not yet received orders. **Note: The red font forms a binary matrix, M with elements $m_{i,j}$.**

The pdf of the inventory levels with stochastic lead times and order crossover (1 of 7)

- A *process* is the sequence of a variable over time ($\{I_t\}, \{W_t\}$ etc).
- A *sub-process* is a subset of the process with identical pipeline states.
- Each sub-process is normally distributed (as the demand is normally distributed and each sub-process is the output of a linear system) and the distribution of the entire process can be multi-modal.
- We require the mean and variance of the inventory levels in each sub-process.
- We obtain this by first determining the distribution of the WIP in each sub-process and then each WIP sub-process is combined with a scaled replenishment order to obtain the **scaled shortfall distribution**.
- A weighted sum of the scaled shortfall distributions in each sub-process then forms the complete inventory distribution.

The pdf of the inventory levels with stochastic lead times and order crossover (2 of 7)

- The inventory in the POUT policy evolves via

$$I_t = T + \mu_D(\mu_L + 1/\beta) - (W_t + O_t/\beta). \tag{12}$$

- For OUT policy (that is, when $\beta = 1$), we can see that the inventory distribution is a reflected **shortfall distribution**, $(W_t + O_t)$, translated by $T + \mu_D(\mu_L + 1/\beta)$ (Zalkind (1978); Robinson et al. (2001)).
- When $\beta \neq 1$ the O_t component is scaled by O_t/β; we call the distribution of $(W_t + O_t/\beta)$ the **scaled shortfall distribution**.
- We require the mean and the variance of the scaled shortfall distribution for each sub-process. The complicating factors are that:
  - O_t is auto-correlated,
  - and the distributions of W_t and O_t/β are correlated with each other.
- As the system is linear, we may exploit the z-transform defined by

$$F(z) = Z\{f[t]\} = \sum_{t=0}^{\infty} f[t]\, z^{-t}. \tag{13}$$

The pdf of the inventory levels with stochastic lead times and order crossover (3 of 7)

- To determine the variance of the WIP in sub-process i, we first note that the variance of the orders maintained by the POUT policy is independent of the lead-time, Disney and Towill (2003), as

$$\frac{\sigma_O^2}{\sigma_D^2} = \sum_{t=0}^{\infty}\left(Z^{-1}\left\{\frac{z\beta}{z+\beta-1}\right\}\right)^2 = \sum_{t=0}^{\infty}\left((1-\beta)^t\beta\right)^2 = \frac{\beta}{2-\beta}, \tag{14}$$

- Here, z is the z-transform operator and

$$Z^{-1}\left\{F(z)\right\} = \frac{1}{2\pi j}\oint_C F(z)\, z^{t-1}\mathrm{d}z = f[t], \tag{15}$$

provides $f[t]$, the inverse z-transform of the transfer function $F(z)$.

- $z\beta(z+\beta-1)^{-1}$ is the transfer function of the orders maintained by the POUT policy under i.i.d. demand with MMSE forecasting, Disney and Towill (2003).
- The relationship between the variance ratio and the sum of the squared impulse response is known as Tsypkin's relationship, Tsypkin (1964), Boute et al. (2022).

Tsypkin's sum of the squared impulse response

Tsypkin's relationship. Patterned on Tsypkin (1964, pp183-192) and Boute et al. (2022)

If the input x_t to a linear system with impulse response function $\tilde{g}_t$ is an i.i.d. random process with the variance $\mathbb{V}[x_t]$, then the variance of output is,

$$\mathbb{V}[y_t] = \mathbb{V}[x_t] \sum_{t=0}^{\infty} (\tilde{g}_t)^2. \quad (16)$$

Proof. Denote $\lim_{t\to\infty} \mathbb{E}[x_t] = \bar{x}$ and $\lim_{t\to\infty} \mathbb{E}[y_t] = \bar{y}$. Taking expectations and limits yields $\bar{y} = \bar{x} \sum_{t=0}^{\infty} \tilde{g}_t$. Indeed, linearity means that a centered input $x_t - \bar{x}$ yields asymptotically centered output $y_t - \bar{y}$. Similarly:

$$\begin{aligned}
\mathbb{V}[y_t] &= \lim_{t\to\infty} \mathbb{E}[(y_t - \bar{y})^2] && \text{(by definition of a variance)} \\
&= \lim_{t\to\infty} \mathbb{E}\left[\left(\sum_{i=0}^{t} (x_i - \bar{x})\tilde{g}_{t-i}\right)^2\right] && \text{(using convolution)} \\
&= \lim_{t\to\infty} \mathbb{E}\left[\left(\sum_{i=0}^{t} (x_i - \bar{x})\tilde{g}_{t-i}\right)\left(\sum_{j=0}^{t} (x_j - \bar{x})\tilde{g}_{t-j}\right)\right] && \text{(expand the square)} \\
&= \lim_{t\to\infty} \sum_{i=0}^{t}\sum_{j=0}^{t} \mathbb{E}[(x_i - \bar{x})(x_j - \bar{x})]\tilde{g}_{t-i}\tilde{g}_{t-j} && \text{(expected value of a sum is the sum of its expected addends)} \\
&= \mathbb{V}[x_t] \sum_{i=0}^{\infty} \tilde{g}_i^2. && (\mathbb{E}[(x_i - \bar{x})(x_j - \bar{x})] = 0 \text{ if } i \neq j \text{ for i.i.d. input } x_t)
\end{aligned}$$

The pdf of the inventory levels with stochastic lead times and order crossover (4 of 7)

- The pdf of the normal distribution with an argument x, mean μ, and standard deviation σ, is defined by

$$\phi\left[x|\mu,\sigma\right] = \frac{e^{-(x-\mu)^2/2\sigma^2}}{\sqrt{2\pi}\sigma}. \tag{17}$$

- Eq. (14) leads to an **uni-modal pdf** for the orders,

$$\phi_O = \phi\left[x|\,\mu_D, \sqrt{\sigma_D^2\beta/(2-\beta)}\right]. \tag{18}$$

- The variance of WIP sub-process i, is given by the variance of the sum of the impulse responses of the open orders,

$$\frac{\sigma_{W,i}^2}{\sigma_D^2} = \sum_{t=0}^{\infty}\sum_{j=1}^{L^+}\left(m_{i,j}Z^{-1}\left\{\frac{\beta z^{1+j}}{z+\beta-1}\right\}\right)^2. \tag{19}$$

where $m_{i,j}$ is an element of the binary matrix **M** that captures whether an order is open or closed.

The pdf of the inventory levels with stochastic lead times and order crossover (5 of 7)

- The distribution of the scaled orders, O_t/β, for all sub-processes is given by

$$\frac{\sigma^2_{O/\beta}}{\sigma^2_D} = \sum_{t=0}^{\infty}\left(Z^{-1}\left\{z/(z+\beta-1)\right\}\right)^2 = \sum_{t=0}^{\infty}\left((1-\beta)^t\right)^2 = \left(2\beta-\beta^2\right)^{-1} = \frac{\sigma^2_O}{\sigma^2_D\beta^2}, \quad (20)$$

which leads to the following pdf,

$$\phi_{O/\beta} = \phi\left[x|\mu_D/\beta, \sqrt{\sigma^2_D/(2\beta-\beta^2}\right]. \quad (21)$$

- The covariance between the WIP and the scaled orders in each sub-process is

$$\operatorname{cov}\left(W_i, O/\beta\right) = \sum_{t=0}^{\infty}\left(Z^{-1}\left\{\sum_{j=1}^{L^+} m_{i,j}\frac{\beta z^{1-j}}{z+\beta-1}\right\}Z^{-1}\left\{\frac{z}{z+\beta-1}\right\}\right) = \operatorname{cov}\left(W_i, O\right)/\beta. \quad (22)$$

The pdf of the inventory levels with stochastic lead times and order crossover (6 of 7) ♠

- $\sigma^2_{I,i}$, the variance of sub-process i in the inventory distribution, is equal to the variance of the shortfall distribution,

$$\sigma^2_{I,i} = \sigma^2_{W,i} + \sigma^2_O/\beta^2 + 2\mathrm{cov}\left(W_i, O\right)/\beta. \tag{23}$$

- The mean of the each of the sub-processes of the inventory distribution is

$$\mu_{I,i} = T + \mu_D\left(\mu_L - \sum\nolimits_{j=1}^{L^+} m_{i,j}\right). \tag{24}$$

- The complete pdf inventory distribution is then given by

$$\phi_I = \sum_{i=1}^{2^{L^+}} q_i\phi\left[x|\,\mu_{I,i}, \sqrt{\sigma^2_D\sigma^2_{I,i}}\right]. \tag{25}$$

- Note, ϕ_I is a multi-modal pdf as it is a combination of the normally distributed pdfs with different means and variances weighted by q_i.
- T, the average inventory, can be set arbitrarily. However, to minimise inventory holding and backlog costs, T is a function of β.

The pdf of the inventory levels with stochastic lead times and order crossover (7 of 7)

- The variance of the complete, multi-modal, inventory is given by

$$\sigma_I^2 = \int_{-\infty}^{\infty} \phi_I(T-x)^2 dx = \sum_{i=1}^{2^{L^+}} q_i \left(\left(\sum_{j=1}^{L^+} m_{i,j} \right)^2 \mu_D^2 + \sigma_D^2 \sigma_{I,i}^2 \right). \quad (26)$$

- The inventory variance contains a weighted sum of the variances of individual sub-processes.
- The mean demand influences the inventory variance; this does not happen with constant lead times.

A numerical example when $L^+ = 4$ ♠

- The table overleaf details the pipeline states ($\mathbf{M}$), the inventory variance, and the mean of each of the $2^{L^+} = 2^4 = 16$ individual inventory sub-processes.
- The variance of each inventory sub-processes is infinite at $\beta = \{0, 2\}$. Furthermore, each sub-process has a single unique minimum, $0 \leq \beta_i^\star < 2$.
- From (11), $\beta_i^\star = 1$ in the sub-processes that do not contain order cross-overs.
- Sub-processes with order-crossover have $\beta_i^\star < 1$ as $\sigma_{W,i}^2 + 2\text{cov}(W_i, O)/\beta$ is convex in β between $\beta = 0$ and $\beta = 1$ and equal to $\sum_{j=1}^{L^+} m_{i,j}$ at $\beta = \{0, 1\}$. This implies there will be a minimum in $\sigma_{I,i}^2$ between $0 \leq \beta < 1$.
- As the complete inventory pdf is a weighted sum of independent variances, some minimised with $\beta = 1$, some minimised with $\beta < 1$, then the β that minimises the variance of the complete inventory distribution is $\beta_\sigma^\star < 1$.
- The proportion of states with order crossover increases in L^+ as the number of states with order crossovers is given by $2^{L^+} - (L^+ + 1)$; β becomes more important as L^+ increases.

Inventory characteristics of each of the sub-processes ♠

M	j				$\sigma^2_{I,i}/\sigma^2_D$	$\mu_{I,i}$	$\beta^\star_i$
i	1	2	3	4			
1	0	0	0	0	$\frac{1}{\beta(2-\beta)}$	$T+\mu\mu_L$	1
2	0	0	0	1	$\frac{-2\beta^5+8\beta^4-12\beta^3+7\beta^2-2\beta-1}{\beta(2-\beta)}$	$T+\mu(\mu_L-1)$	0.656633
3	0	0	1	0	$\frac{2\beta^4-6\beta^3+5\beta^2-2\beta-1}{\beta(2-\beta)}$	$T+\mu(\mu_L-1)$	0.689845
4	0	0	1	1	$\frac{2\beta^5-10\beta^4+16\beta^3-10\beta^2+4\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-2)$	0.60974
5	0	1	0	0	$\frac{2\beta^3-3\beta^2+2\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-1)$	0.751274
6	0	1	0	1	$\frac{2\beta^5-6\beta^4+10\beta^3-8\beta^2+4\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-2)$	0.676129
7	0	1	1	0	$\frac{-2\beta^4+6\beta^3-6\beta^2+4\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-2)$	0.689845
8	0	1	1	1	$\frac{-2\beta^5+8\beta^4-12\beta^3+9\beta^2-6\beta-1}{\beta(2-\beta)}$	$T+\mu(\mu_L-3)$	0.656633
9	1	0	0	0	$\frac{-\beta^2+\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-1)$	1
10	1	0	0	1	$\frac{-2\beta^4+6\beta^3-6\beta^2+4\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-2)$	0.689845
11	1	0	1	0	$\frac{2\beta^3-4\beta^2+4\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-2)$	0.751274
12	1	0	1	1	$\frac{2\beta^4-6\beta^3+7\beta^2-6\beta-1}{\beta(2-\beta)}$	$T+\mu(\mu_L-3)$	0.689845
13	1	1	0	0	$\frac{-2\beta^2+4\beta-1}{\beta(2-\beta)}$	$T+\mu(\mu_L-2)$	1
14	1	1	0	1	$\frac{2\beta^3-5\beta^2+6\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-3)$	0.751274
15	1	1	1	0	$\frac{-3\beta^2+6\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-3)$	1
16	1	1	1	1	$\frac{-4\beta^2+8\beta+1}{\beta(2-\beta)}$	$T+\mu(\mu_L-4)$	1
Overall	-	-	-	-	-	T	$\beta^\star_\sigma=0.73$

Probability of the pipeline being in state i, q_i

- The previous table details the first- and second-order moments of each of the sub-processes for a maximum lead time of $L^+ = 4$. These hold regardless of the lead-time probabilities.
- To identify the probability of the pipeline being in state i, we assume $\left\{p_0 = \frac{1}{2}, p_1 = p_2 = p_3 = 0, p_4 = \frac{1}{2}\right\}$
  - This extreme set of probabilities exaggerates the impact of order crossovers.
  - Representative of a setting where 50% of orders are sent via ship with $L = 4$ and 50% are sent by air with $L = 0$. The air shipments arrive before the next order is made.
- The maximum lead time is $L^+ = 4$ and the average lead time is $\mu_L = 2$.
- The probability the pipeline is in state i is

$$q_i = \frac{1}{2^{L^+}} \prod_{j=1}^{L^+} \left[1 + (-1)^{\upsilon} \left(2 \sum_{k=j}^{L^+} p_k - 1 \right) \right]. \tag{27}$$

- For our setting here, the probability the WIP pipeline is in state i, is $\forall i, q_i = 0.0625$.

The inventory distribution maintained by the OUT policy for 90% availability ($\beta = 1$) ♠

Using the information in the sub-process table, and that $\forall i, q_i = 0.0625$ in (25) we can construct the following figure of the inventory pdf.

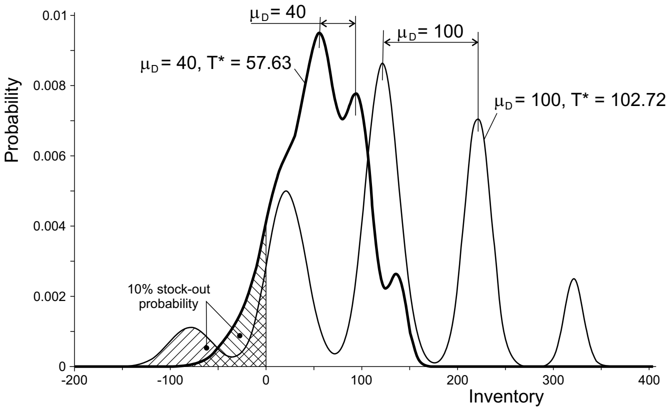

There are five modes in the inventory pdf as the 16 WIP states have only five unique means, $\mu_{I,i}$; each mode is separated by μ_D.

The OUT policy is not optimal with stochastic lead times

- Consider linear convex inventory holding (h) and backlog positions (b) costs exist

$$J = h\mathbb{E}[[I_t]^+] + b\mathbb{E}[[-I_t]^+]. \tag{28}$$

- The derivative at $\beta = 1$ is given by

$$\left.\frac{\mathrm{d}J}{\mathrm{d}\beta}\right|_{\beta=1} = \frac{\sigma(b+h)e^{-(T+\mu^2)/4\sigma^2}}{64\sqrt{\pi}} \times \left(6+3\sqrt{2}\right)e^{-(T^2+6T\mu+\mu^2)/8\sigma^2} + 4\sqrt{6}e^{-(T^2+6T\mu+3\mu^2)/12\sigma^2}\Bigg) > 0, \tag{29}$$

  implying that, for all cost combinations, the OUT policy is never optimal as there always exists a $\beta < 1$ that is more economical.

- When $\mu_D = 100$, the inventory levels have a variance of 10,300 for the OUT policy. Numerical experiments reveal a single POUT policy **minimises inventory variance at** $\beta^\star_\sigma = 0.73$ of 10,280–0.2% less than the OUT variance.

- For the $\mu_D = 40$ case, the inventory variance maintained by the OUT policy is 1900. The numerically optimized POUT feedback parameter is the same, $\beta^\star_\sigma = 0.73$, producing an inventory variance of 1879–a 1% reduction.

Even for only inventory costs, the POUT policy outperforms the OUT policy when stochastic lead time are present.

After setting $\{\beta^\star, T^\star\}$ optimally, the figure describes the percentage economic gain $\left((J_{OUT} - J_{POUT})\, J_{OUT}^{-1} \times 100\%\right)$, from using the POUT policy.

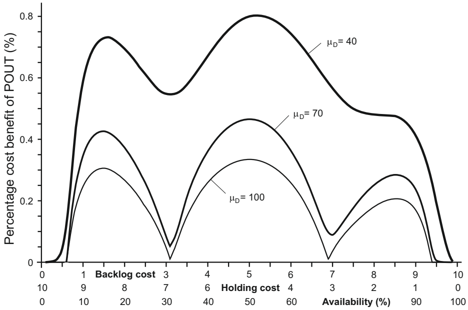

The inventory variance consists of two components. The first component is due to the modes which depend on μ_D but are independent of β. The second component is a function of the variance of each sub-process which depends on β. The modes are dominant; the inventory cost benefit of the POUT policy is limited and decreases in μ_D.

The inventory cost optimal $\beta^\star$ in the POUT policy

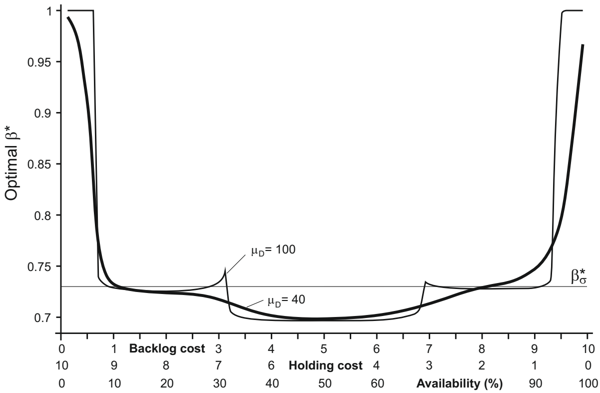

- $\beta^\star$ is near unity when the availability target is (very) near 0% or 100%, but for most availability targets $\beta^\star \approx 0.725$.
- Almost always, $\beta^\star \neq \beta^\star_\sigma$ implying that the tightest inventory control does not always lead to the minimal cost.

Safety stock requirements with inventory cost optimal $\beta^\star$

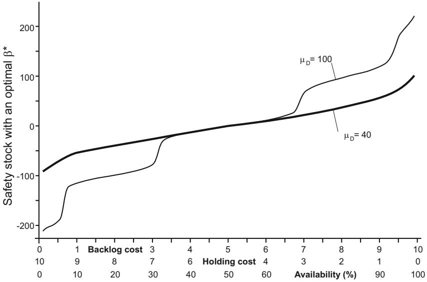

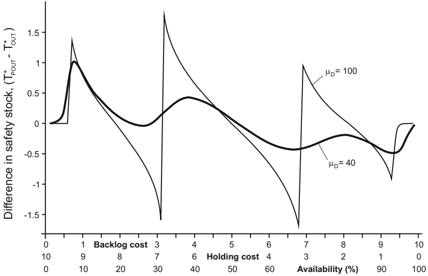

Left panel. Safety stock requirements when $\beta^\star$ is used for different cost ratios. The multi-modal nature of the $\mu_D = 100$ case results in rapid increases in the safety stock requirement at predictable availability % related to the multi-modal inventory pdf.

Right panel. The difference, $T_{POUT} - T_{OUT}$, reveals minimum costs do not always concur with minimal safety stock. This is contrary to the constant lead time case where the least cost solution always has the smallest safety stocks.

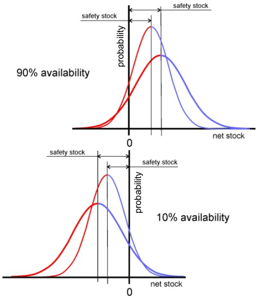

Constant lead time

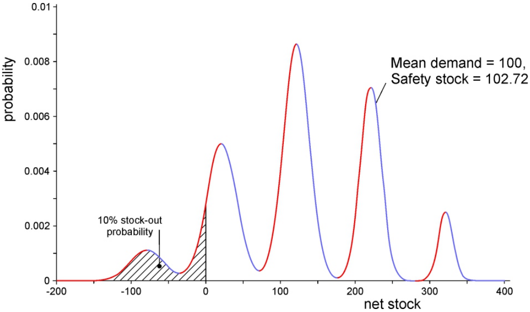

Stochastic lead time

Consider a constant lead time where the cost ratio is such that the optimal availability is less than 50%. To minimise the safety stock requirements, $\beta = 0$ leads to $\sigma_I = \infty$ and a safety stock requirement of $T^\star = -\infty$. However, inventory costs $J = \infty$. Under a stochastic lead time, due to the multi-modality of the inventory distribution, these *trailing edge* effects can happen in the whole availability range.

Bullwhip ratio achieved when an inventory cost minimising $\beta^\star$ is used in the POUT policy ♠

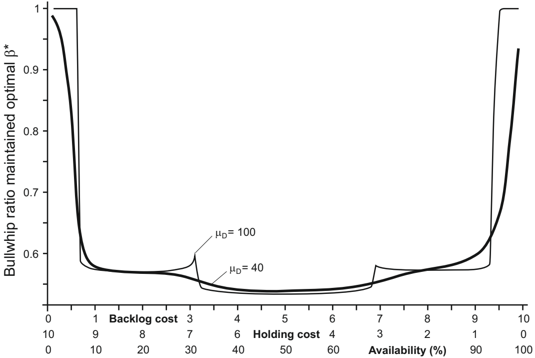

A 40% reduction in order variance is possible between 8% to 92% availability despite the objective function consisting only of inventory-related costs. If costs are associated with order variability are also present, these will also be reduced. This is an important result as order variance costs are somewhat harder to quantify, Lee et al. (2000).

Summary: OUT and POUT performance when the lead time distribution is known beforehand

Constant lead time	Stochastic lead time
Optimal availability is the critical fractile ($b/(b+h)$)	
Minimizing inventory variance minimizes inventory costs	Minimizing inventory costs almost never occurs with minimal inventory variance
Minimizing safety stocks minimizes inventory costs	Minimizing safety stocks does not result in minimal inventory costs
OUT policy inventory cost optimal	POUT policy is always more economic than the OUT policy (even for inventory costs alone)
Bullwhip ratio is equal to unity when an inventory cost optimal OUT policy is used	An inventory cost optimal POUT policy almost always results in a 20% reduction of the bullwhip effect

Note: 20% reduction in bullwhip is the average for the 12 empirical lead time distribution between China and the US.

Quoted and actual lead times in a global shipping lane

Sometimes, we don't know the distribution of the lead times; they must be forecasted.

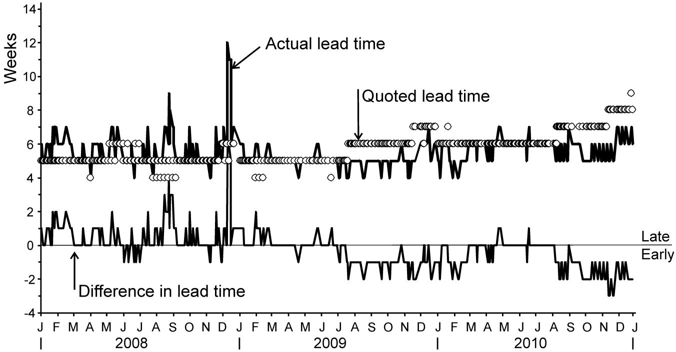

Michna et al. (2020) study the consequences of forecasting AR(1) demand and i.i.d. lead times with the moving average forecasting mechanism.

The OUT policy when lead times have to be forecasted, Michna et al. (2020)

- Suppose we have AR(1) demand,

$$D_t = \mu_D + \rho(D_{t-1} - \mu_D) + \epsilon_t, \tag{30}$$

  that we forecast with an n period moving average,

$$\hat{D}_t = \frac{1}{n}\sum_{i=1}^{n} D_{t-i}. \tag{31}$$

- We forecast the lead times from the previously observed deliveries using

$$\hat{L}_t = \frac{1}{m}\sum_{i=1}^{m} L_{t-i-L^+}\ , \tag{32}$$

  Note L_{t-i-L^+} are lead times guaranteed to have been observed by the manufacturer at the beginning of a period $t-i$.

- The lead time demand (the OUT level, S_t is then given by

$$\hat{D}_t^L = \hat{L}_t\hat{D}_t = \frac{1}{mn}\sum_{i=1}^{n} D_{t-i}\left(\sum_{i=1}^{m} L_{t-i-L^+}\right). \tag{33}$$

- The replenishment orders are then given by

$$O_t = \hat{D}_t^L - \hat{D}_{t-1}^L + D_{t-1}. \tag{34}$$

When the lead time has been forecasted the order distribution becomes multi-modal ♠

- When the AR(1) demand and i.i.d. lead times have been forecasted with the moving average (with demand averaged over n periods and lead times averaged over m periods), the orders and the inventory distribution become multi-modal.

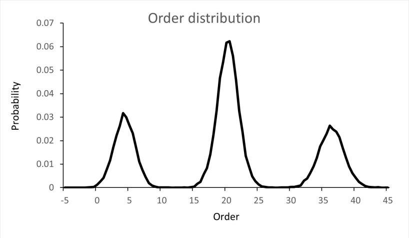

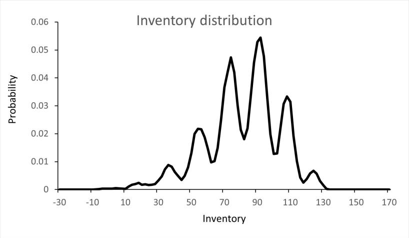

- Note: Here $\sigma_\epsilon^2 = 1$, $\mu_D = 20$, $\rho = 0.5$, $m = 5$, $n = 5$, $p_0 = 0.5$ and $p_4 = 0.5$.
- Here, the order distribution has 3 modes, and inventory distribution has 8 modes.
- When the lead time distribution is known beforehand, the order distribution is uni-modal and the inventory distribution has 5 modes when $p_0 = 0.5$ and $p_4 = 0.5$.

Forecasting lead times is a major cause of the Bullwhip effect ♠

- The following bullwhip ratio can be found,

$$BM = \frac{2\sigma_L^2}{n^2 m^2}\left(m(1-\rho^n) + \frac{n(1+\rho)}{1-\rho} - \frac{(1+\rho^2)(1-\rho^n)}{(1-\rho)^2}\right) + \frac{2\sigma_L^2 \mu_D^2}{\sigma_D^2 m^2} + \left(\frac{2\mu_L^2}{n^2} + \frac{2\mu_L}{n}\right)(1-\rho^n) + 1. \quad (35)$$

- Here, the red font denotes the bullwhip generated by the OUT policy with moving average forecasts of AR(1) demand with constant lead times.
- The mean demand influences the bullwhip ratio when the demand has to be forecasted; this does not happen under constant lead times.
- This bullwhip expression can be found using the **law of total variance**,

$$\mathbb{V}[O_t] = \mathbb{E}[\mathbb{V}[O_t|L]] + \mathbb{V}[\mathbb{E}[O_t|L]]. \quad (36)$$

Details omitted for brevity, see Michna et al. (2020).

Bullwhip characterisation when moving average is used the forecast AR(1) demand and i.i.d. lead times

$$\lim_{m\to\infty} BM = 1 + (1-\rho^n)\left(\frac{2\mu_L^2}{n^2} + \frac{2\mu_L}{n}\right). \tag{37}$$

$$BM_{n=1} = \frac{\rho 2\sigma_D^4\left(\sigma_L^2 - m\left(m\mu_L(\mu_L+1)+\sigma_L^2\right)\right)}{m^2} + \frac{2\mu_D^2\sigma_D^2\sigma_L^2 + m\sigma_D^4\left(2m\mu_L(\mu_L+1)+2\sigma_L^2+m\right)}{m^2}. \tag{38}$$

$$BM_{iid} = 1 + \frac{2\mu_D^2\sigma_L^2}{m^2\sigma_D^2} + \frac{2\sigma_L^2(m+n-1)}{m^2n^2} + \frac{2\mu_L(\mu_L+n)}{n^2}. \tag{39}$$

$$BM_{\rho\to1} = 1 + \frac{2\sigma_L^2\left(\mu_D^2+\sigma_D^2\right)}{m^2\sigma_D^2}. \tag{40}$$

$$BM_{\rho\to-1} = 1 + \frac{2\mu_D^2\sigma_L^2}{m^2\sigma_D^2} - \frac{(2m-1)\left((-1)^n-1\right)\sigma_L^2}{m^2n^2} - \frac{2\left((-1)^n-1\right)\mu_L(\mu_L+n)}{n^2} \tag{41}$$

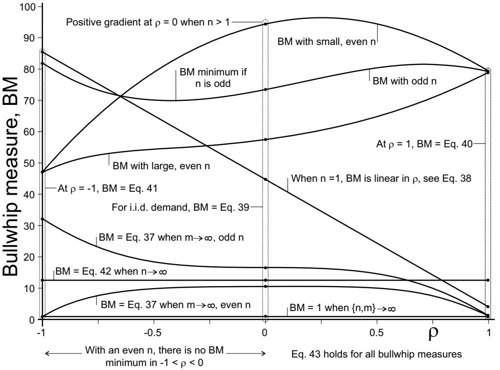

$$\lim_{n\to\infty} BM = 1 + \frac{2\mu_D^2\sigma_L^2}{m^2\sigma_D^2}. \tag{42}$$

$$BM = \frac{2\sigma_L^2}{n^2m^2}\left(m(1-\rho^n) + \frac{n(1+\rho)}{1-\rho} - \frac{(1+\rho^2)(1-\rho^n)}{(1-\rho)^2}\right) + \frac{2\sigma_L^2\mu_D^2}{\sigma_D^2 m^2} + \left(\frac{2\mu_L^2}{n^2} + \frac{2\mu_L}{n}\right)(1-\rho^n) + 1. \tag{43}$$

Many real lead times are auto-correlated

- Material 95047502 from a Danish manufacturing company has 1349 lead time observations, Disney et al. (2024).
- It's time series and its auto-correlation function are presented below. The region between the dotted lines illustrates a 95% confidence interval for zero correlation.
- The Akaike information criterion function provides an auto-regressive moving average model ARMA(1,2) with coefficients $\phi_1 = 0.6426$, $\theta_1 = -0.1822$ and $\theta_2 = 0.0718$ (a stationary time series ARMA(1,2) model) with a mean $\mu_D = 5.5690$.

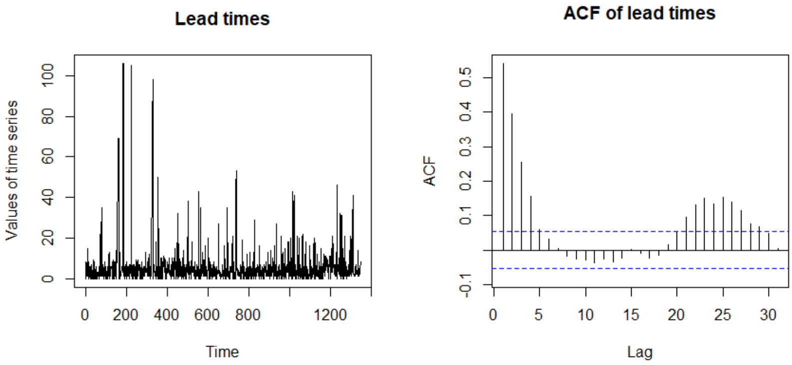

Theorem. Under AR(1) demand and auto-correlated lead times the bullwhip ratio is given by

$$\frac{\mathbb{V}[O_t]}{\mathbb{V}[d_t]} = \frac{1-\rho^2 - 2\rho(1-\rho^2)\mathbb{E}[\rho^{\widehat{L}_t}] + 2\rho^2\mathbb{E}[\rho^{2\widehat{L}_t}] - 2\rho^3\mathbb{E}[\rho^{\widehat{L}_t+\widehat{L}_{t-1}}]}{(1-\rho)^2} + \frac{2\mu_D^2}{\sigma_D^2}\mathbb{V}[\widehat{L}_t]\left(1-\mathbf{Corr}[\widehat{L}_t, \widehat{L}_{t-1}]\right). \quad (44)$$

Proof. First, let us define

$$L = \{\widehat{L}_{t-1}, \widehat{L}_t\}. \quad (45)$$

We then evoke the **law of total variance** to find the variance of O_t,

$$\mathbb{V}[O_t] = \mathbb{V}[\mathbb{E}[O_t|L]] + \mathbb{E}[\mathbb{V}[O_t|L]]. \quad (46)$$

Using (34), the variance of O_t conditional upon L is given by

$$\mathbb{V}[O_t|L] = \mathbb{V}\left[\left\{D_{t-1} + \sum_{k=0}^{\widehat{L}_t-1}(\mu_D + \rho^{k+1}(D_{t-1}-\mu_D)) - \sum_{k=0}^{\widehat{L}_{t-1}-1}(\mu_D + \rho^{k+1}(D_{t-2}-\mu_D))\right\}|L\right]. \quad (47)$$

Proof of Theorem, continued

Incorporating the previous demand (the first addend) into the sum in the second addend provides

$$\mathbb{V}[O_t|L] = \mathbb{V}\left[\left\{\sum_{k=-1}^{\widehat{L}_t-1}(\mu_D + \rho^{k+1}(D_{t-1} - \mu_D)) - \sum_{k=0}^{\widehat{L}_{t-1}-1}(\mu_D + \rho^{k+1}(D_{t-2} - \mu_D))\right\}|L\right]. \quad (48)$$

Ignoring the mean values (as they don't influence the variance) allows the following simplification,

$$\mathbb{V}[O_t|L] = \mathbb{V}\left[\left\{\sum_{k=-1}^{\widehat{L}_t-1}\rho^{k+1}D_{t-1} - \sum_{k=0}^{\widehat{L}_{t-1}-1}\rho^{k+1}D_{t-2}\right\}|L\right]. \quad (49)$$

Closing the sums,

$$\mathbb{V}[O_t|L] = \mathbb{V}\left[\left\{\frac{1-\rho^{\widehat{L}_t+1}}{1-\rho}D_{t-1} - \rho\frac{1-\rho^{\widehat{L}_{t-1}}}{1-\rho}D_{t-2}\right\}|L\right]. \quad (50)$$

The variance of a sum is the sum of the variance of the addends minus twice the covariance between the addends,

$$\begin{aligned}\mathbb{V}[O_t|L] =& \mathbb{V}\left[\left\{\frac{1-\rho^{\widehat{L}_t+1}}{1-\rho}D_{t-1}\right\}|L\right] + \mathbb{V}\left[\left\{\rho\frac{1-\rho^{\widehat{L}_{t-1}}}{1-\rho}D_{t-2}\right\}|L\right] \\ &- 2\mathbf{Cov}\left[\left\{\frac{1-\rho^{\widehat{L}_t+1}}{1-\rho}D_{t-1}, \rho\frac{1-\rho^{\widehat{L}_{t-1}}}{1-\rho}D_{t-2}\right\}|L\right]. \quad (51)\end{aligned}$$

Proof of Theorem, continued

Taking the variances and the covariance provides,

$$\mathbb{V}[O_t|L] = \sigma_D^2\left(\left(\frac{1-\rho^{\widehat{L}_t+1}}{1-\rho}\right)^2 + \left(\rho\frac{1-\rho^{\widehat{L}_{t-1}}}{1-\rho}\right)^2\right) - 2\rho\sigma_D^2\left(\frac{1-\rho^{\widehat{L}_t+1}}{1-\rho}\right)\left(\rho\frac{1-\rho^{\widehat{L}_{t-1}}}{1-\rho}\right), \quad (52)$$

which by stationarity of $\widehat{L}_t$ gives

$$\mathbb{E}[\mathbb{V}[O_t|L]] = \frac{\mathbb{E}[1-\rho^{1+\widehat{L}_t}]^2 + \rho^2\mathbb{E}[1-\rho^{\widehat{L}_t}]^2 - 2\rho^2\mathbb{E}[(1-\rho^{1+\widehat{L}_t})(1-\rho^{\widehat{L}_{t-1}})]}{(1-\rho)^2}\sigma_D^2. \quad (53)$$

We can compute the second term of (46), the variance of O_t, as follows,

$$\mathbb{V}[\mathbb{E}[O_t|L]] = \mathbb{V}\left[\mathbb{E}\left[\left\{\sum_{k=-1}^{\widehat{L}_t-1}(\mu_D+\rho^{k+1}(D_{t-1}-\mu_D)) - \sum_{k=0}^{\widehat{L}_{t-1}-1}(\mu_D+\rho^{k+1}(D_{t-2}-\mu_D))\right\}|L\right]\right]$$
$$= \mathbb{V}\left[\sum_{k=-1}^{\widehat{L}_t-1}\mu_D - \sum_{k=0}^{\widehat{L}_{t-1}-1}\mu_D\right] = \mathbb{V}[(\widehat{L}_t+1)\mu_D - \widehat{L}_{t-1}\mu_D] = 2\mu_D^2(\mathbb{V}[\widehat{L}_t] - \mathbf{Cov}[\widehat{L}_t,\widehat{L}_{t-1}]). \quad (54)$$

Combining (54) with (53) provides (44). ∎

Correlated two point distribution for lead times with naïve forecasts

- Consider a stationary sequence of lead times $\{L_t\}_{t=-\infty}^{\infty}$ with $\mathbf{Corr}[L_t, L_{t-1}] = \rho_L$ with a two point marginal distribution; $\mathbb{P}[L_t = a] = 1/2$, $\mathbb{P}[L_t = b] = 1/2$, where a, b, $a \neq b$ are positive integers.
- Assume the joint distribution of (L_{t-1}, L_t) is

	$L_t = a$	$L_t = b$
$L_{t-1} = a$	$\frac{1}{4} + \frac{1}{4}\rho_L$	$\frac{1}{4} - \frac{1}{4}\rho_L$
$L_{t-1} = b$	$\frac{1}{4} - \frac{1}{4}\rho_L$	$\frac{1}{4} + \frac{1}{4}\rho_L$

- Note; $\mathbb{E}[L_t] = (a+b)/2$, $\mathbb{V}[L_t] = (a-b)^2/4$, $\mathbf{Corr}[L_t, L_{t-1}] = \rho_L$, $\mathbb{E}[\rho^{L_t}] = 1/2\rho^a + 1/2\rho^b$, $\mathbb{E}[\rho^{2L_t}] = 1/2\rho^{2a} + 1/2\rho^{2b}$ and $\mathbb{E}[\rho^{L_t + L_{t-1}}] = \rho^{2a}(1/4 + \rho_L/4) + \rho^{a+b}(1/2 - \rho_L/2) + \rho^{2b}(1/4 + \rho_L/4)$.
- The bullwhip measure for naïve forecasting of a two point lead time distribution is

$$\left.\frac{\mathbb{V}[O_t]}{\mathbb{V}[D_t]}\right|_{\text{Naïve}} = \frac{1 - \rho(1-\rho^2)(\rho^a + \rho^b) + \rho^2(\rho^{2a} + \rho^{2b} - 1) - \frac{1}{2}\rho^3(2\rho^{a+b}(1-\rho_L) + (\rho^{2a} + \rho^{2b})(1+\rho_L))}{(1-\rho)^2} + \frac{\mu_D^2 (a-b)^2}{2\sigma_D^2}(1-\rho_L)\ . \quad (55)$$

Correlated two point distribution for lead times with naïve forecasts: The case of $a = 2$ and $b = 4$

The bullwhip ratio with the naïve forecasting of lead times as a function of demand correlation $-1 \leq \rho \leq 1$ and lead time correlation $-1 \leq \rho_L \leq 1$ for $\mu_D = 10$ and $\sigma_D = 2$ is shown below.

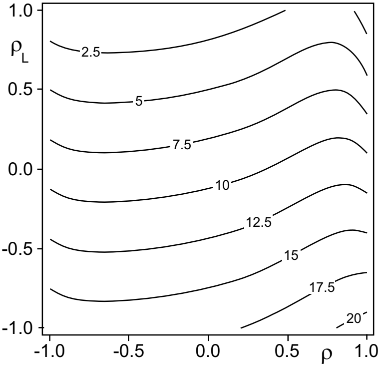

Note, the bullwhip ratio is linear and decreasing in ρ_L.

Concluding remarks

- The stochastic lead time problem is a vast, rich, and under explored research stream.
- The nature of the impact of the stochastic lead time depends strongly on the assumptions made.
- A wide variety of mathematical tools and methods are required.
- There are plenty of open research questions in this area, including:
    - Identifying optimal replenishment policies.
    - Understanding the impact of more complex/realistic demand patterns.
    - Investigating the impact of cross correlation between auto-correlated demand and auto-correlated lead times.
    - State dependent lead times, Boute et al. (2014).
    - Seasonal lead times.
    - More complex supply chain settings such as dual-sourcing and vendor managed inventory.

Thank you for listening

Inventory control with stochastic lead times: A tutorial

Stephen M. Disney

Bibliography I

Balakrishnan, A., J. Geunes, M.S. Pangburn. 2004. Coordinating supply chains by controlling upstream variability propagation. *Manufacturing & Service Operations Management* **6**(2) 163–183.

Boute, R.N., S.M. Disney, J. Gijsbrechts, J.A. Van Mieghem. 2022. Dual sourcing and smoothing under nonstationary demand time series: Reshoring with SpeedFactories. *Management Science* **68**(2) 1039–1057.

Boute, R.N., S.M. Disney, M.R. Lambrecht, B. Van Houdt. 2014. Coordinating lead times and safety stocks under autocorrelated demand. *European Journal of Operational Research* **232**(1) 52–63.

Boute, R.N., J.A. Van Mieghem. 2015. Global dual sourcing and order smoothing: The impact of capacity and lead times. *Management Science* **61**(9) 2080–2099.

Disney, S.M., A. Maltz, X. Wang, R.D.H. Warburton. 2016. Inventory management for stochastic lead times with order crossovers. *European Journal of Operational Research* **248** 473–486.

Disney, S.M., Z. Michna, P. Nielsen. 2024. The bullwhip effect with correlated lead times and auto-correlated demand. *Working paper* 27pp.

Disney, S.M., D.R. Towill. 2003. On the bullwhip and inventory variance produced by an ordering policy. *Omega* **31**(3) 157–167.

Eppen, G.D., R.K. Martin. 1988. Determining safety stock in the presence of stochastic lead time and demand. *Management Science* **34**(11) 1380–1390.

Feller, W. 1958. *An Introduction to Probability Theory and Its Applications*. New York: John Wiley & Sons, Inc., Vol. 1, 268-276.

Lee, H.L., K.C. So, C.S. Tang. 2000. The value of information sharing in a two-level supply chain. *Management Science* **46**(5) 626–643.

Michna, Z., S.M. Disney, P. Nielsen. 2020. The impact of stochastic lead times on the bullwhip effect under correlated demand and moving average forecasts. *Omega* **93** 102033.

Naim, M.M. 1993. *Learning curve models for predicting the performance of industrial systems*. PhD thesis, University of Wales College of Cardiff, United Kingdom.

Robinson, L.W., J.R. Bradley, L.J. Thomas. 2001. Consequences of order crossover under order-up-to inventory policies. *Manufacturing & Service Operations Management* **3**(3) 175–188.

Tsypkin, Y.Z. 1964. *Sampling Systems Theory and its Application, Vol. 2*. Pergamon Press, Oxford.

Zalkind, D. 1978. Order-level inventory systems with independent stochastic leadtimes. *Management Science* **24**(13) 1384–1392.